

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:

Hyun-kwon CHUNG, et al.

Serial No. 09/903,630

Group Art Unit: 2145

Confirmation No. 1050

Filed: July 13, 2001

Examiner: Jeffrey R. Swearingen

For: REPRODUCING APPARATUS AND SERVER SYSTEM PROVIDING ADDITIONAL INFORMATION THEREFOR

APPEAL BRIEF UNDER 37 C.F.R § 41.37

Mail Stop Appeal Brief-Patents
Commissioner for Patents
PO Box 1450
Alexandria, VA 22313-1450

Sir:

Pursuant to the Appellant's earlier filed Notice of Appeal on October 23, 2006, Appellant hereby appeals to the Board of Patent Appeals and Interferences from the final rejection mailed June 22, 2006. Appellant submits this Appeal Brief along with the filing fee of \$500.00 set forth in 37 C.F.R. §41.20(b)(2).

Also enclosed is a Claims Appendix in compliance with 37 C.F.R. § 41.37(c)(1)(viii) and an Evidence Appendix in compliance with 37 C.F.R. § 41.37(c)(1)(ix) includes references of record and discussed in the Appeal Brief below. A Related Proceedings Appendix in compliance with 37 C.F.R. § 41.37(c)(1)(x) is enclosed and indicated as being NONE.

I. Real Party in Interest

Pursuant to 37 C.F.R. §41.37(c)(1)(i), due to the assignment executed on July 27, 2001 by the inventors Hyun-kwon CHUNG and Jung-kwon HEO and recorded in the United States Patent and Trademark Office at Reel 012262, Frame 0320, the real party in interest is as follows:

Samsung Electronics Co., Ltd.
416, Maetan-dong, Paldal-gu,
Suwon-city, Kyungki-do
Republic of Korea

II. Related Appeals and Interferences

Pursuant to 37 C.F.R. §41.37(c)(1)(ii), although the real party in interest has other appeals and interferences, there is only one other pending appeal and interference believed to directly affect or be directly affected by, or have any bearing upon the decision of the Board of Patent Appeals and Interferences in this appeal. The related appeal is for U.S. patent application no. 10/995,295. No decision on this appeal has been made at this time.

III. Status of Claims

Pursuant to 37 C.F.R. §41.37(c)(1)(iii), claims 45 through 51, 53 through 60, and 62 through 71 are pending in this application at the filing of this Appeal Brief. Claims 45 through 51, 53 through 60, and 62 through 71 stand finally rejected. Claims 45, 51, and 59 are independent claims, and claims 46 through 50, 53 through 58, 60, and 62 through 71 are dependent claims.

In view of the final Office Action mailed June 22, 2006 as supplemented in the Advisory Action of October 17, 2006, claims 45 through 51, 53 through 60, and 62 through 71 stand finally rejected. This Appeal Brief is an appeal of the finally rejected claims 45 through 51, 53 through 60, and 62 through 71.

IV. Status of Amendments

Pursuant to 37 C.F.R. §41.37(c)(1)(iv), no amendments have been filed since the final Office Action of June 22, 2006. Pursuant to 37 C.F.R. §41.37(c)(viii), a copy of the claims involved in the appeal is included in their present condition is included in the Claims Appendix.

V. Summary of the Claimed Subject Matter

Pursuant to 37 C.F.R. §41.37(c)(v) and without admitting to the applicability of 35 U.S.C. §112, paragraph 6 for any claim element in these claims, non-limiting examples from the specification have been incorporated to explain aspects of the invention in the independent claims.

According to an aspect of the invention of claim 45, a reproduction apparatus for reproducing contents according to independent claim 45 includes an identifier provider for providing an identifier of the contents (Paragraphs 0021, 0022 and FIGs. 1 and 4 describe a recording apparatus 10 that reads an identifier, such as an International Standard Recording Code (ISRC), from an optical disc 1 in operation 401); a network connector (Paragraph 0022 and FIG. 1 describe a network connector 13); and a controller for storing the contents identifier provided by the identifier provider as a Cookie file (Paragraphs 0026 and 0030 and FIGs. 1 and 4 describe the an identification generator 11 which transmits the identifier to a controller 12, and a browser 14 stores the transmitted identifier in a Cookie file in operation 402), transmitting the stored contents identifier through the network connector to a server system, which provides additional information related to the contents through the network connector (Paragraphs 0030, 0031 and FIG. 4 describe that the browser 14 is called and connects to the server system 100 through a network connector 13 in operation 405, and once connected, the browser 14 retrieves the stored Cookie file from a memory, and provides the Cookie file including the identifier to the server 102 in operation 406), and receiving through the network connector the additional information provided from the server system after the stored contents identifier was transmitted (Paragraphs 0024-0026, 0031 and FIG. 4 describe that the browser 14, receives through the network connector 13, the additional information from the server 102 in operation 407, the additional information being determined by the server 100 to be associated with the contents according to the identifier in the transmitted Cookie file).

According to an aspect of the invention of claim 51, a server system providing additional information items includes an additional information database which stores additional information items related to a plurality of contents (Paragraphs 0004, 0026 to 0028 and FIGs. 1 and 2 describe an additional information database 101 including information related to the contents using an ISRC); and a server for receiving a file including an identifier of predetermined contents from a reproduction apparatus for reproducing the contents (Paragraphs 0030, 0031,

0033 and FIGs. 1, 4 and 5 describe that the server 102 connects to the browser 14 through a network connector 13 of a reading apparatus 10 and in operation 502, the server 102 requests the identifier from the browser 14 and receives the Cookie file with the identifier), the file being prepared by and stored by a browser on the reproduction apparatus prior to transmission to the server (paragraph 0026, 0030, 0031 and FIGs. 1, 4, and 5 describe that the identifier is stored by the browser 14 in a Cookie file in operation 402, and the Cookie file is transmitted in operation 406), retrieving one of the additional information items related to the contents identifier from the additional information data base according to the received file (Paragraph 0033 and FIGs. 1, 4 and 5 describe that the server 102 extracts the additional information from an additional information database 101 according to the identifier in the received the Cookie file in operation 503), and transmitting the retrieved one additional information item to the reproduction apparatus (Paragraph 0033 and FIGs. 1, 4 and 5 describe that the server 102 transmits the additional information to the browser 14 in operation 504), where the contents identifier is recorded in at least one recording medium on which the contents are recorded (Paragraph 0022 and FIGs. 1 and 4 describe that the identifier is read from an optical disc 1 in operation 401).

According to an aspect of the invention of claim 59, an information storage medium for use with a recording and/or reproducing apparatus and having a Cookie program which implements a method of generating a Cookie file used by the apparatus, the method comprising: detecting an identifier of predetermined contents (Paragraphs 0004, 0021, 0022, 0029, 0030 and FIGs. 1 and 4 describe a reproducing apparatus 10 that reads and detects an identifier, such as an International Standard Recording Code (ISRC), from an optical disc 1 in operation 401, the contents being audio contents or video contents, and the additional information can be words of songs, personal information items of singers, recent activities, and/or other like songs); preparing and storing the detected contents identifier in the Cookie file for use in a subsequent transmission by the apparatus to a server system providing additional information related to the predetermined contents through a network connector of the apparatus in response to the sent

Cookie file (Paragraphs 0030, 0031; FIGs. 1 and 4 describe an identification generator 11 that transmits the identifier to a controller 12, and a browser 14 stores the transmitted identifier in a Cookie file in operation 402, the browser 14 sends the Cookie file through the network connector 13 to a server system 100 in operation 405, and the server system 100 provides the additional information based upon the identifier in the Cookie file in operations 406 and 407), where the contents identifier is an international standard recording code (ISRC) read from a recording medium (Paragraphs 0021, 0022, 0030 and FIGs. 1 and 4 describe the identification generator 11 that transmits the identifier such as an ISRC read from a disc 1).

VI. Grounds of Rejection to be Reviewed

As per 37 CFR 41.37(c)(1)(vi), the following is a concise statement of each ground of appeal.

1. Whether claims 45-51, 53-60, and 62-71 are patentable under 35 U.S.C. §103 in view of Meyer et al. (U.S. Patent No. 6,829,368) and Montulli (U.S. Patent No. 5,744,670).

VII. Argument

1. **Claims 45-51, 53-60, and 62-71 are patentably distinguishable over Meyer et al. and Montulli**

In general, in order to reject a claim under 35 U.S.C. §103, a combination of references must be provided which discloses each element of the claim in the manner recited in the claim. In interpreting the reference, the Examiner is to broadly interpret the claim, but must do so within the bounds of reason. In re Morris, 127 F.3d 1048, 1053-55;44 U.S.P.Q.2d 1023, 1027-28 (Fed. Cir. 1997). Thus, while the Examiner is to avoid reading limitations from the specification into the claims, the Examiner should not interpret claim limitations so broadly as to contradict or otherwise render a limitation meaningless as would be understood by those of ordinary skill in the art. See, In re Cortright, 165 F.3d 1353, 1357-58; 49 U.S.P.Q.2d 1464, 1467 (Fed. Cir. 1999), In re Zletz, 893 F.2d 319, 321-22; 13 U.S.P.Q.2d 1320, 1322 (Fed. Cir. 1989).

Additionally, where the Examiner is relying on a feature as being inherently disclosed in a reference, it is incumbent on the Examiner to provide evidence that such a feature both

necessarily exists in the reference, and exists in a manner which is the same as presented in the claims. As noted by the Federal Circuit in In re Robertson, 169 F.3d 743, 745; 49 U.S.P.Q.2d 1949, 1950-51 (Fed. Cir. 1999), to "establish inherency, the extrinsic evidence 'must make clear that the missing descriptive matter is necessarily present in the thing described in the reference, and that it would be so recognized by persons of ordinary skill.' *Continental Can Co. v. Monsanto Co.* , 948 F.2d 1264, 1268, 20 U.S.P.Q.2d 1746, 1749 (Fed. Cir. 1991). 'Inherency, however, may not be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient.'" *Id.* at 1269, 20 U.S.P.Q.2d at 1749 (quoting In re Oelrich, 666 F.2d 578, 581, 212 U.S.P.Q. 323, 326 (C.C.P.A. 1981)."

Further, in order to establish a prima facie obviousness rejection, the Examiner needs to provide both the existence of individual elements corresponding to the recited limitations, and a motivation to combine the individual elements in order to create the recited invention. The Examiner is further required to evaluate the record as a whole, and to account for contrary teachings existing in the record. In re Young, 927 F.2d 588; 18 U.S.P.Q.2d 1089 (Fed. Cir. 1991) cited by MPEP 2143.01. Should the Examiner fail to provide evidence that either one of the individual elements or the motivation does not exist in the prior art, then the Examiner has not provided sufficient evidence to maintain a prima facie obviousness rejection of the claim. In re Kotzab, 217 F.3d 1365; 55 U.S.P.Q.2d 1313 (Fed. Cir. 2000). Thus, the burden is initially on the Examiner to provide particular evidence as to why one of ordinary skill in the art would have been motivated to combine the individual elements to create the recited invention, and to demonstrate that this evidence existed in the prior art. In re Zurko, 258 F.3d 1379; 59 U.S.P.Q.2d 1693 (Fed. Cir. 2001).

In view of the above and as set forth below, there is insufficient evidence of a record that Meyer et al. (U.S. Patent No. 6,829,368) and Montulli (U.S. Patent No. 5,744,670) discloses explicitly or on intently, the features of the claims in a manner supporting a prima facie obviousness rejection under 35 U.S.C. §103.

A. The combination of Meyer et al. and Montulli does not disclose storing a contents identifier prior to transmission through a network connector as recited in claims 45-50, and 65-69

By way of review, claim 45 recites, among other features, "a network connector," and "a controller for storing the contents identifier provided by the identifier provider as a Cookie file, transmitting the stored contents identifier through the network connector to a server system, which provides additional information related to the contents through the network connector, and receiving through the network connector the additional information provided from the server system after the stored contents identifier was transmitted."

In contrast, Meyer et al. discloses a decoder which collects identifiers in response to a user request while the objects containing these identifiers are being played. In order to capture the identifier, the decoder includes an interface having a button used by the user to request information about the objects. In the example of FIG. 2, an MP3 ripper extracts metadata from a read CD, uses the metadata to obtain information about songs of the CD from a database Cddb, and uses this obtained information to embed an object identifier (OID) in the resulting MP3 audio file. When selected, the decoding device packages a message including the identifier read from the audio file, and invokes a communication application, such as an Internet Browser. The invoked communication application forwards the provided identifier as a message to a server. (Col. 4, lines 13-31, col. 5, lines 3-6, col. 6, lines 34-59, col. 13, lines 4-27, col. 16, line 61 to col. 52; FIGs. 1 and 2).

However, while the communication application forwards the identifier in a message prepared by the decoder apparatus, there is no suggestion that the communication application stores the content identifier as opposed to the decoder interface or the decoder apparatus, or that the ISRC or identifier should be stored in a file such as a Cookie file. As such, to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a Cookie file.

In order to cure this deficiency, the Examiner asserts on page 2 of the Final Office Action mailed June 22, 2006 that storage inherently precedes transmission and provides a portion Kurose et al., "Computer Networking: A Top Down Approach Featuring the Internet," pp. 383-385 (Addison Wesley Longman Inc. (2001)) evidencing that a network interface card includes RAM. As an initial point of clarification, it is noted that the instant application has a U.S. filing date of July 13, 2001. In contrast, the copyright date of Kurose et al. is 2001, indicating that Kurose et al. was published sometime in 2001. As such, there is insufficient evidence of record as to whether Kurose et al. was published before or after July 13, 2001 such that the record is unclear to the extent to which Kurose et al. represents the state of the art necessarily utilized by the Meyer et al.

Moreover, to the extent Kurose et al. represents a state of the art, Kurose et al. suggests that a network interface card (NIC) has RAM for local storage with the NIC. Kurose et al. further explains that such NICs are semi-autonomous as compared to the parent node, even though the NIC may be in the same box as the node (i.e., computer). (Pgs. 383-384 of Kurose et al.) As such, Kurose et al. stands for the proposition that, when a data frame is received from the node for transmission, the RAM of the NIC is used in the transmission of the received data frame apart from the node and is operating under its own controller. As such, any inherent storage prior to transmission occurring in Kurose et al. occurs in the semi-autonomous NIC and does not necessarily occur in the node. Thus, under principles of inherency, there remains insufficient evidence that storage occurs outside of a network interface card and prior to transmission through the network interface card, or that a controller of the node could control such storage within the RAM of the semi-autonomous NIC.

Since the Examiner has not provided such evidence corresponding to the invention as claimed, it is respectfully submitted that the Examiner has not provided sufficient evidence to rely on Meyer et al. inherently disclosing storage of the ISRC in the decoder as opposed to an network interface card which semi-autonomously operates to perform a transmission as set forth

in the Office Action and as is required to disclose the features of claim 45.

Since Montulli is not relied upon as disclosing this feature, it is respectfully submitted that the combination does not disclose or suggest the invention as recited in claim 45.

Claims 46-50 and 65-69 are patentable due at least to their depending from claim 45.

B. The combination of Meyer et al. and Montulli does not suggest performing local storage as a Cookie file prior to transmission as recited in claims 45-50 and 65-69

Additionally, as Meyer et al. does not disclose transmitting a contents identifier in a Cookie file, the Examiner relies upon Montulli to disclose that it is well known to use Cookie files for data transmission. As evidence of a motivation, the Examiner asserts that one skilled in the art would have been motivated to use Cookie files for data transmission since Cookie files were known in the art as a data transmission scheme utilized in web browsers and are a type of container which could be used to transmit both the content identifier of a type described in Montulli as well as the state information actually described in Montulli. The essence of the Examiner's argument is that, since cookie technology was well known in the art, one skilled in the art could have utilized the cookie to transmit contents identifiers instead of state information.

However, an unsubstantiated statement that existing elements could be combined as it was in the skill of the art to do so does not provide a basis for a rejection under 35 U.S.C. 103(a). In re Fine, 837 F2d 1071; 5 U.S.P.Q.2d 1596 (Fed. Cir. 1988). Similarly, an unsubstantiated statement that elements could be combined as being "common sense" does not provide a basis for a rejection under 35 U.S.C. §103(a) since such unsupported statements prevent meaningful review under the Administrative Procedures Act, 5 U.S.C. §706. In re Zurko, 258 F.3d 1379; 59 U.S.P.Q.2d 1693 (Fed. Cir. 2001). In essence, there needs to be proof that such a motivation exists, not conjecture. This rigorous proof is required in order to prevent the trap of impermissible hindsight.

In view of the above, to the extent Meyer et al. suggests the decoder preparing a message and invoking the Internet browser solely to transmit the contents identifier in a

message, Meyer et al. does not suggest a need for the creation of the Cookie file with the contents identifier since Meyer et al. does not rely upon the Internet browser for more than relaying messages. Thus, to the extent that the message of Meyer et al. is a container for the contents identifier, the record is unclear as to why one skilled in the art would utilize cookie technology as the particular delivery mechanism for this container.

Moreover, Montulli teaches that such use of Cookie files is restricted to the context of client-server relationships where the state of the client in this relationship cannot otherwise be maintained in client-server system. (Col. 2, lines 15-21 of Montulli). Thus, Montulli discloses using Cookie files sent from a server to a client so that, when the client later accesses the same server, the client information is maintained with respect to the server by sending the cookie to the same server. (Col. 9, lines 45-65 and Example 1 of Montulli). There is no suggestion that the same or another Cookie file is pre-generated at the client prior to contact with the server so as to be independent of the Cookie file provided by the server.

There is further no suggestion of an advantage to using the Cookie files in such other contexts, and there is no suggestion that the message in Meyer et al. is needed to maintain a client-server relationship such that the contents identifier of Meyer et al. would not be understood as a type of state information described in the context of Montulli. As such, any such wide use relied upon by the Examiner has not been shown with respect to other contexts, such as that in Meyer et al. Thus, even assuming that cookie technology is in widespread use to maintain a client-server relationship, there is insufficient evidence that such wide spread use, in and of itself, is sufficient to evidence why one skilled in the art would be motivated to modify Meyer et al., which does not describe messaging in the context of maintenance of a client-server relationship, to use cookie technology as a message/container for a contents identifier.

Moreover, to the extent that the Cookie file can be used to coordinate clients and servers, Meyer et al. teaches that the coordination between the client and server is taken into account using sets of rules which govern timing of returned data after the identifier has been

sent to the server. (Col. 5, lines 48-53). There is no suggestion that a Cookie file like that of Montulli represents an improvement over this set of rules already suggested in Meyer et al.

Also, while Meyer et al. discloses a decoder which collects identifiers, the identifiers can be an ISRC and can be read from a media object. The decoding device packages a message including the identifier, and invokes a communication application, such as an Internet Browser. The invoked communication application forwards the provided identifier as a message to a server 1. (Col. 3, lines 54-59, col. 13, lines 4-27, col. 16, line 61 to col. 52). When received at the server 1, the server 1 uses the identifier to return data or actions associated with the received identifier. Additionally, context data stored by the decoder may be used at the server to determine information on the user and/or the computer so as to judge, for instance whether the returned data or actions need to be altered to be age appropriate. (Col. 5, lines 25-47). However, while Meyer et al. discloses forwarding a message including the read ISRC and the context information, there is no suggestion that the ISRC or identifier should be stored in a file such as a Cookie file. As such, to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a Cookie file as suggested in Montulli.

In reviewing the combined teachings, Montulli discloses the use of the Cookie file in the context of client-server relationships, where the state of the client cannot otherwise be maintained in a client-server system. (Col. 2, lines 15-21 of Montulli). Thus, Montulli discloses a server creating the Cookie file and sending the Cookie file from the server to the client. Therefore, when the client later accesses the same server, the client information is maintained with respect to the server by the client sending the Cookie file to the same server. (Col. 9, lines 45-65 and Example 1 of Montulli). There is no suggestion that the same or another Cookie file is pre-generated at the client prior to contact with the server so as to be independent of the Cookie file provided by the server.

In contrast, Meyer et al. suggests storing the identifiers in the object itself, which is locally stored under the control of the decoder/MP3 player/receiver, and when a connection to a server 1 is needed, the decoder/MP3 player/receiver extracts the identifier and sends the identifier via the Internet Browser to the server 1. (Col. 6, line 60 to col. 7, line 5; FIGs. 1 and 2). There is no suggestion of a need to again store the already-stored identifier in this process.

Additionally, while suggested as being useful for retaining client status information in the client-server system after the server has disconnected from the client, there is no suggestion that the Cookie file of Montulli should be used by the client to transmit non-transmission related data (such as contents identifiers) not generated during a prior or existing network connection between the client and the server. Specifically, in Meyer et al., the decoder/MP3 player/receiver sends the identifier prior to establishing a client-server relationship as the client, whereas Montulli suggests that the Cookie file is generated by a server to enable the server to maintain an existing client-server relationship. (Col. 2, lines 13-21). Montulli does not suggest any particular advantage for the use of Cookie files in other contexts, such as where the client is sending stored information to the server. Since the prior art as a whole does not suggest using the Cookie file outside of maintaining a client-server relationship, the record as a whole does not provide sufficient evidence as to why one skilled in the art would further modify the decoder of Meyer et al. to implement the teaching of Montulli to have a client store and send the Cookie file incorporating an ISRC or identifier.

Indeed, there are concerns about the over use of Cookie files as suggested in Lin et al., *Taking a Byte Out of Cookies: Privacy, Consent, and the Web*, ACM Policy Proceedings of the Ethics and Social Impact Component on Shaping Policy in the Information Age, pp. 39-51 (1998). At the very least, Lin et al. would limit the extent to which one skilled in the art would have wanted to implement the Cookie file unless needed to maintain a client-server relationship. Thus, to the extent Cookie files are suggested by Montulli, the record as a whole also cautions against overuse of Cookie files. Therefore, in view of Lin et al., one skilled in the art would not

be motivated to use Cookie files every time messages are to be sent to a server from a client instead of merely to maintain the server-client relationship, which is the suggestion actually made by Montulli.

As such, the record is unclear as to why one skilled in the art would modify Meyer et al., which transfers messages generated by the decoder apparatus through the Internet browser, to instead use the Internet browser to store, locally, the Cookie file of Montulli, and then send the stored Cookie file with the same information as the message to the server. Thus, it is respectfully submitted that, in view of the record as a whole, there is insufficient of a motivation to combine the technologies of Meyer et al. and Montulli to store the message of Meyer et al. in a Cookie file as is required to sustain a *prima facie* obviousness rejection of claims 45-50 and 65-69.

C. The combination of Meyer et al. and Montulli does not disclose a browser storing a contents identifier prior to transmission as recited in claims 51, 53-57 and 70

By way of review, claim 51 recites, among other features, "a server for receiving a file including an identifier of predetermined contents from a reproduction apparatus for reproducing the contents, the file being prepared by and stored by a browser on the reproduction apparatus prior to transmission to the server."

In contrast and as similarly noted above in Section VII(1)(A), Meyer et al. teaches invoking the communication application to forward the provided identifier as a message to a server. (Col. 4, lines 13-31, col. 5, lines 3-6, col. 6, lines 34-59, col. 13, lines 4-27, col. 16, line 61 to col. 52; FIGs. 1 and 2). However, while the communication application forwards the identifier in a message prepared by the decoder apparatus, there is no suggestion that the communication application stores the content identifier as opposed to the decoder interface or the decoder apparatus. Instead, Meyer et al. discloses a decoder/MP3 player/receiver which obtains, embeds, and/or retrieves the content identifier, with the internet browser being used at most to forward messages under the control of the decoder/MP3 player/receiver. As such, to

the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a locally stored file, such as the Cookie file as suggested in Montulli.

Further, to the extent that Kurose et al. suggests that network interface cards (NICs) have RAM, there is no suggestion that the communication application necessarily controls the NIC to store the ISRC in the RAM of the NIC. Instead, Kurose et al. teaches that such NICs are semi-autonomous as compared to the parent node such that any storage occurs independent of the existence of the communication application, even though the NIC may be in the same box as the node (i.e., computer). (Pgs. 383-384 of Kurose et al.) Thus, any inherent storage prior to transmission occurring in Kurose et al. occurs in the semi-autonomous NIC and does not necessarily occur in the node. Thus, under principles of inherency there remains insufficient evidence that storage occurs outside of a network interface card and prior to transmission through the network interface card, or that a browser of the node could control such storage within the RAM of the semi-autonomous NIC.

Since Montulli is not relied upon as disclosing this feature, it is respectfully submitted that the combination does not disclose or suggest the invention as recited in claim 51.

Claims 53-57 and 70 are patentable due at least to their depending from claim 51.

D. The combination of Meyer et al. and Montulli does not suggest the use of a browser to perform local storage of contents identifier prior to transmission as recited in claims 51, 53-58, and 70

As noted above in Section VII(1)(B) and VII(1)(C), Meyer et al. discloses a decoder which collects identifiers. The identifiers can be an ISRC and can be read from a media object. The decoding device packages a message including the identifier, and invokes a communication application, such as an Internet browser. The invoked communication application forwards the provided identifier as a message to a server 1. (Col. 3, lines 54-59, col. 13, lines 4-27, col. 16, line 61 to col. 52). When received at the server 1, the server 1 uses the identifier to return data

or actions associated with the received identifier. Additionally, context data stored by the decoder may be used at the server to determine information on the user and/or the computer so as to judge, for instance whether the returned data or actions need to be altered to be age appropriate. (Col. 5, lines 25-47). However, while Meyer et al. discloses forwarding a message including the read ISRC and the context information, there is no suggestion that the ISRC or identifier should be locally stored in a file. As such, to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a file, such as the Cookie file as suggested in Montulli.

Further, while suggested as being useful for retaining client status information in the client-server system, there is no suggestion that the Cookie file of Montulli would be useful in the context of the reading and transmission of the ISRC in Meyer et al. or otherwise resolve a problem in Meyer et al. As also noted above in Section VII(1)(B), since the prior art as a whole does not suggest using the Cookie file outside of maintaining a client-server relationship and teaches away from overuse of the Cookie file, the record as a whole does not provide sufficient evidence as to why one skilled in the art would further modify the decoder of Meyer et al. to implement the teaching of Montulli to have a client store and send the Cookie file incorporating an ISRC or identifier.

As such, it is respectfully submitted that there is insufficient evidence of record as to why one skilled in the art would utilize the Cookie file of Montulli in the context of Meyer et al. in a manner meeting the features of claim 51 as is required to maintain a prima facie obviousness rejection under 35 U.S.C. §103.

Similarly, there is insufficient evidence of record to maintain a prima facie obviousness rejection under 35 U.S.C. §103 as to claims 53-58, and 70.

E. The combination of Meyer et al. and Montulli does not disclose a browser storing a contents identifier as a file prior to transmission through a

network connector and a server to receive the stored file as recited in claims 54-56

By way of review, claim 54 recites, among other features, "a network connector," and "a controller to control the browser to prepare and store the file including the identifier from the identifier provider prior to transmission to the server, to use the browser to transmit the stored identifier to the server through the network connector, to use the browser to receive the retrieved one additional information item transmitted from the server through the network connector corresponding to the transmitted identifier, and to control a display of the received one additional information item."

As similarly set forth in Sections VII(1)(A) and (1)(C), while the communication application of Meyer et al. forwards the identifier in a message prepared by the decoder apparatus, there is no suggestion that the communication application stores the content identifier as opposed to the decoder interface or the decoder apparatus, or that the ISRC or identifier should be stored in a file such as a Cookie file. As such, to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a Cookie file as suggested in Montulli. Instead, Meyer et al. discloses a decoder/MP3 player/receiver which obtains, embeds, and/or retrieves the content identifier, with the internet browser being used at most to forward messages under the control of the decoder/MP3 player/receiver.

To the extent Kurose et al. represents the state of the art, Kurose et al. relates to the network interface cards (NICs) having RAM, which are separate from the node (i.e., the decoder of Meyer et al.). Kurose et al. further explains that such NICs are semi-autonomous as compared to the parent node, even though the NIC may be in the same box as the node (i.e., computer). (Pgs. 383-384 of Kurose et al.) As such, Kurose et al. stands for the proposition that, when a data frame is received from the node for transmission, the RAM of the NIC is used in the transmission of the received data frame apart from the node and is operating under its

own controller. As such, any inherent storage prior to transmission occurring in Kurose et al. occurs in the semi-anonymous NIC and does not necessarily occur in the node. Thus, under principles of inherency there remains insufficient evidence that storage occurs outside of a network interface card and prior to transmission through the network interface card, or that a controller of the node could control such storage within the RAM of the semi-anonymous NIC.

Since the Examiner has not provided such evidence corresponding to the invention as claimed, it is respectfully submitted that the Examiner has not provided sufficient evidence to rely on Meyer et al. inherently disclosing such features as set forth in the Office Action and as is required to disclose the features of claim 54.

Since Montulli is not relied upon as disclosing this feature, it is respectfully submitted that the combination does not disclose or suggest the invention as recited in claim 54.

Claims 55 and 56 are deemed patentable due at least to the patentability of claim 54.

F. The combination of Meyer et al. and Montulli does not suggest the use of a browser to perform local storage of contents identifier as a Cookie File prior to transmission where same contents identifier is on information storage medium being read by apparatus as recited in claim 58

By way of review, claim 58, which depends from claim 57 recites, among other features, the that "the server receives the contents identifier from the browser installed in the controller of the reproduction apparatus as a Cookie file prepared by the browser." As similarly noted above in Sections VII(1)(B) and VIII(1)(D), while suggested as being useful for retaining client status information in the client-server system, there is no suggestion that the Cookie of Montulli would be useful in the context of the reading and transmission of the ISRC in Meyer et al. or otherwise resolve a problem in Meyer et al. As also noted above in Section VII(1)(B), since the prior art as a whole does not suggest using the Cookie file outside of maintaining a client-server relationship and teaches away from overuse of the Cookie file, the record as a whole does not provide sufficient evidence as to why one skilled in the art would further modify the communication application of the decoder of Meyer et al. to implement the teaching of Montulli to have a client store and send the Cookie file incorporating an ISRC or identifier prior to transmission. As such,

it is respectfully submitted that there is insufficient evidence of record as to why one skilled in the art would utilize the Cookie file of Montulli in the context of Meyer et al. in a manner meeting the features of claim 58 as is required to maintain a prima facie obviousness rejection under 35 U.S.C. §103.

G. The combination of Meyer et al. and Montulli does not disclose an information storage medium having a Cookie program for implementing a method including storing as a Cookie File a contents identifier prior to transmission through a network connector as recited in claims 59, 60, 62-64, and 71

By way of review, claim 59 recites, among other features, "a Cookie program which implements a method of generating a Cookie file used by the apparatus" including "preparing and storing the detected contents identifier in the Cookie file for use in a subsequent transmission by the apparatus to a server system providing additional information related to the predetermined contents through a network connector of the apparatus in response to the sent Cookie file."

In contrast and as similarly noted in Section VII(A), to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a Cookie file as suggested in Montulli. Instead, Meyer et al. discloses a decoder/MP3 player/receiver which obtains, embeds, and/or retrieves the content identifier, with the internet browser being used at most to forward messages under the control of the decoder/MP3 player/receiver.

Moreover, to the extent Kurose et al. represents a state of the art, Kurose et al. stands for the proposition that, when a data frame is received from the node for transmission, the RAM of the NIC is used in the transmission of the received data frame apart from the node and is operating under its own controller. As such, any inherent storage prior to transmission occurring in Kurose et al. occurs in the semi-autonomous NIC and does not necessarily occur in the node. Thus, under principles of inherency there remains insufficient evidence that storage occurs

outside of a network interface card and prior to transmission through the network interface card, or that a controller of the node could control such storage within the RAM of the semi-anonymous NIC.

Since the Examiner has not provided such evidence corresponding to the invention as claimed, it is respectfully submitted that the Examiner has not provided sufficient evidence to rely on Meyer et al., inherently disclosing such features as set forth in the Office Action and as is required to disclose the features of claim 59.

Since Montulli is not relied upon as disclosing this feature, it is respectfully submitted that the combination does not disclose or suggest the invention as recited in claim 59.

Claims 60, 62-64, and 71 are patentable due at least to their depending from claim 59.

H. The combination of Meyer et al. and Montulli does not suggest an information storage medium having a Cookie program for implementing a method including preparing and storing the detected contents identifier in the Cookie file for use in a subsequent transmission by the apparatus as recited in claims 59, 60, 62-64, and 71

As similarly noted above in Sections VII(1)(B) and VII(1)(F), while Meyer et al. discloses forwarding a message including the read ISRC and the context information, there is no suggestion that the ISRC or identifier should be locally stored in a file. As such, to the extent Meyer et al. teaches reading an ISRC, temporarily buffering the ISRC, and including the ISRC in a message as set forth in col. 13, lines 24-51, there is no suggestion of a need or benefit to again locally store the ISRC in a file, such as the Cookie file as suggested in Montulli.

Further, while suggested as being useful for retaining client status information in the client-server system, there is no suggestion that the Cookie file of Montulli would be useful in the context of the reading and transmission of the ISRC in Meyer et al. or otherwise resolve a problem in Meyer et al. As also noted above in Section VII(1)(B), since the prior art as a whole does not suggest using the Cookie file outside of maintaining a client-server relationship and teaches away from overuse of the Cookie file, the record as a whole does not provide sufficient evidence as to why one skilled in the art would further modify the decoder of Meyer et al. to

implement the teaching of Montulli to have a client store the Cookie file incorporating an ISRC or identifier.

As such, it is respectfully submitted that there is insufficient evidence of record as to why one skilled in the art would utilize the Cookie file of Montulli in the context of Meyer et al. in a manner meeting the features of claim 59 as is required to maintain a prima facie obviousness rejection under 35 U.S.C. §103.

For similar reasons, it is respectfully submitted that insufficient evidence of record to maintain a prima facie obviousness rejection under 35 U.S.C. §103 of claims 60, 62-64, and 71.

I. The combination of Meyer et al. and Montulli does not suggest the use of a browser to perform local storage of contents identifier as a Cookie File prior to transmission through a network connector as recited in claim 66

By way of review, claim 66, which depends from claim 45, recites, among other features, "a browser," where "the controller controls the browser to prepare and store the Cookie file with the contents identifier in the prepared Cookie file, and uses the browser to transmit the stored Cookie to the server system and to receive the additional information provided from the server system after transmitting the stored Cookie."

In contrast and as similarly noted above in Sections VII(1)(B) and (1)(F), while suggested as being useful for retaining client status information in the client-server system, there is no suggestion that the Cookie file of Montulli would be useful in the context of the reading and transmission of the ISRC in Meyer et al. or otherwise resolve a problem in Meyer et al. As also noted above in Section VII(1)(B), since the prior art as a whole does not suggest using the Cookie file outside of maintaining a client-server relationship and teaches away from overuse of the Cookie file, the record as a whole does not provide sufficient evidence as to why one skilled in the art would further modify the communication application of the decoder of Meyer et al. to implement the teaching of Montulli to have a client store and send the Cookie file incorporating an ISRC or identifier prior to transmission. As such, it is respectfully submitted that there is insufficient evidence of record as to why one skilled in the art would utilize the Cookie file of

Montulli in the context of Meyer et al. in a manner meeting the features of claim 66 as is required to maintain a prima facie obviousness rejection under 35 U.S.C. §103.

VIII. Conclusion

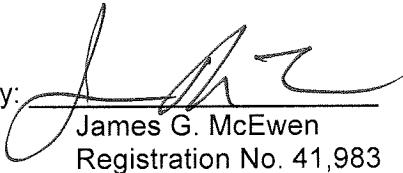
In view of the law and facts stated herein, the Appellant respectfully submits that the Examiner has failed to cite a reference or combination of references sufficient to maintain an obviousness rejection of the rejected claims and has failed to rebut the arguments in at least the Amendment After Final Rejection filed September 20, 2006 and/or the Amendment filed April 4, 2006.

For all the foregoing reasons, the Appellant respectfully submits that the cited prior art does not teach or suggest the presently claimed invention. The claims are patentable over the prior art of record and the Examiner's findings of unpatentability regarding claims 45 through 51, 53 through 60 and 62 through 71 should be reversed.

The Commissioner is hereby authorized to charge any additional fees required in connection with the filing of the Appeal Brief to our Deposit Account No. 50-3333.

Respectfully submitted,

STEIN, MCEWEN & BUI LLP

By: 
James G. McEwen
Registration No. 41,983

1400 Eye Street, NW, Suite 300
Washington, D.C. 20005
Telephone: (202) 216-9505
Facsimile: (202) 216-9510

Date: March 20, 2007

IX. Claims Appendix

1-44. (CANCELED)

45. (PREVIOUSLY PRESENTED) A reproduction apparatus for reproducing contents, comprising:

an identifier provider for providing an identifier of the contents;
a network connector; and
a controller for storing the contents identifier provided by the identifier provider as a Cookie file, transmitting the stored contents identifier through the network connector to a server system, which provides additional information related to the contents through the network connector, and receiving through the network connector the additional information provided from the server system after the stored contents identifier was transmitted.

46. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, further comprising a reading unit for reading data from at least one storage medium, in which the contents are stored, and reads the contents identifier from the at least one storage medium, wherein the identifier provider provides the read contents identifier read from the at least one storage medium to the controller.

47. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, further comprising a reading unit for reading data from at least one storage medium, in which the contents are stored, and reads an international standard recording code (ISRC) from the at least one storage medium, wherein the identifier provider receives the ISRC read and provides the ISRC as the contents identifier to the controller.

48. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, further comprising:

a reading unit for reading the contents from at least one storage medium in which the contents are stored; and
a reproducer for reproducing contents read by the reading unit.

49. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 48, wherein the reproducer further comprises a decoder for decoding the read contents.

50. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 49, wherein the reproducer further comprises:

a speaker for receiving audio data output from the decoder and delivering sound; and
a display apparatus for receiving video data output from the decoder and displaying images.

51. (PREVIOUSLY PRESENTED) A server system providing additional information items, comprising:

an additional information database which stores additional information items related to a plurality of contents; and

a server for receiving a file including an identifier of predetermined contents from a reproduction apparatus for reproducing the contents, the file being prepared by and stored by a browser on the reproduction apparatus prior to transmission to the server, retrieving one of the additional information items related to the contents identifier from the additional information data base according to the received file, and transmitting the retrieved one additional information item to the reproduction apparatus,

wherein the contents identifier is recorded in at least one recording medium on which the contents are recorded.

52. (CANCELED)

53. (PREVIOUSLY PRESENTED) The server system of claim 51, wherein:

an international standard recording code (ISRC) is recorded in at least one recording medium on which the contents are recorded, and

the server receives the ISRC code as the contents identifier, retrieves the one of the additional information items mapped to the ISRC code from the additional information database, and transmits the retrieved one additional information item to the reproduction apparatus.

54. (PREVIOUSLY PRESENTED) The server system of claim 51, wherein:

the server transmits the additional information item corresponding to the received contents identifier to the reproduction apparatus, and

the reproduction apparatus comprises:

an identifier provider for providing the identifier of the contents,
the browser,

a network connector, and
a controller to control the browser to prepare and store the file including the identifier from the identifier provider prior to transmission to the server, to use the browser to transmit the stored identifier to the server through the network connector, to use the browser to receive the retrieved one additional information item transmitted from the server through the network connector corresponding to the transmitted identifier, and to control a display of the received one additional information item.

55. (PREVIOUSLY PRESENTED) The server system of claim 54, wherein:
the reproduction apparatus further comprises a reading unit for reading data from at least one storage medium,
the at least one storage medium stores the contents,
the identifier provider provides the contents identifier read from the at least one storage medium to the controller, and
the controller receives the contents identifier from the reproduction apparatus for transmitting the contents identifier provided by the identifier provider through the network connector to the server.

56. (PREVIOUSLY PRESENTED) The server system of claim 55, wherein the server receives an international standard recording code (ISRC) read from the at least one storage medium by the reading unit and provides the received ISRC code as the contents identifier to the controller.

57. (PREVIOUSLY PRESENTED) The server system of claim 51, wherein the server receives the contents identifier from the browser installed in a controller of the reproduction apparatus.

58. (PREVIOUSLY PRESENTED) The server system of claim 57, wherein the server receives the contents identifier from the browser installed in the controller of the reproduction apparatus as a Cookie file prepared by the browser.

59. (PREVIOUSLY PRESENTED) An information storage medium for use with a recording and/or reproducing apparatus and comprising a Cookie program which implements a method of generating a Cookie file used by the apparatus, the method comprising:

detecting an identifier of predetermined contents; and
preparing and storing the detected contents identifier in the Cookie file for use in a subsequent transmission by the apparatus to a server system providing additional information related to the predetermined contents through a network connector of the apparatus in response to the sent Cookie file,

wherein the contents identifier is an international standard recording code (ISRC) read from a recording medium.

60. (PREVIOUSLY PRESENTED) The information storage medium of claim 59, wherein the method further comprises reading the contents identifier from the recording medium on which the contents are stored.

61. (CANCELED)

62. (PREVIOUSLY PRESENTED) The information storage medium of claim 60, wherein the Cookie file is prepared by the apparatus and is stored on the apparatus prior to transmission, and the Cookie file includes the contents identifier read from the recording medium.

63. (PREVIOUSLY PRESENTED) The information storage medium of claim 60, wherein the Cookie file is prepared by a browser provided in the apparatus and is stored prior to transmission, and the Cookie file includes the contents identifier read from the recording medium.

64. (PREVIOUSLY PRESENTED) The information storage medium of claim 60, wherein the Cookie file is prepared by a browser provided in the apparatus and is stored prior to transmission, and the method further comprises transmitting the Cookie file to the server system providing the additional information through a network.

65. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, wherein the received additional information is reproduced without reproducing the corresponding contents.

66. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, further comprising a browser, and the controller controls the browser to prepare and store the Cookie file with the contents identifier in the prepared Cookie file, and uses the browser to transmit the

stored Cookie to the server system and to receive the additional information provided from the server system after transmitting the stored Cookie.

67. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, wherein:
the controller receives an input requesting retrieval of the additional information,
if the received input requests receipt of the additional information without reproducing the corresponding contents, the additional information is retrieved from the server system using the Cookie file without reproducing the corresponding contents, and
if the received input requests the additional information while reproducing the corresponding contents, the additional information is retrieved from the server system using the Cookie file while reproducing the corresponding contents.

68. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, further comprising a memory in which the controller stores the Cookie file prior to providing the Cookie file to the server system.

69. (PREVIOUSLY PRESENTED) The reproduction apparatus of claim 45, wherein the contents comprises audio and/or video predetermined contents, and the additional information includes words of a song of the audio and/or video contents, personal information items on singers of the audio and/or video contents, contents of recent activities of the audio and/or video contents, other songs of a similar genre of the audio and/or video contents, or combinations thereof.

70. (PREVIOUSLY PRESENTED) The server system of claim 51, wherein the contents comprises audio and/or video predetermined contents, and the additional information items include words of a song of the audio and/or video contents, personal information items on singers of the audio and/or video contents, contents of recent activities of the audio and/or video contents, other songs of a similar genre of the audio and/or video contents, or combinations thereof.

71. (PREVIOUSLY PRESENTED) The information storage medium of claim 59, wherein the predetermined contents comprises audio and/or video predetermined contents, and the additional information includes words of a song of the audio and/or video contents, personal information items on singers of the audio and/or video contents, contents of recent activities of

the audio and/or video contents, other songs of a similar genre of the audio and/or video contents, or combinations thereof.

X. Evidence Appendix

1. Meyer et al. (U.S. Patent No. 6,829,368)
2. Montulli (U.S. Patent No. 5,744,670)
3. Kurose et al., "Computer Networking: A Top Down Approach Featuring the Internet," pp. 383-385 (Addison Wesley Longman Inc. (2001))
4. Lin et al., *Taking a Byte Out of Cookies: Privacy, Consent, and the Web*, ACM Policy Proceedings of the Ethics and Social Impact Component on Shaping Policy in the Information Age, pp. 39-51 (1998)



US006829368B2

(12) **United States Patent**
Meyer et al.

(10) **Patent No.:** US 6,829,368 B2
(45) **Date of Patent:** Dec. 7, 2004

(54) **ESTABLISHING AND INTERACTING WITH ON-LINE MEDIA COLLECTIONS USING IDENTIFIERS IN MEDIA SIGNALS**

(75) Inventors: **Joel R. Meyer**, Portland, OR (US); **Geoffrey B. Rhoads**, West Linn, OR (US)

(73) Assignee: **Digimarc Corporation**, Tualatin, OR (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 583 days.

(21) Appl. No.: 09/769,017

(22) Filed: Jan. 24, 2001

(65) **Prior Publication Data**

US 2001/0031066 A1 Oct. 18, 2001

Related U.S. Application Data

(63) Continuation-in-part of application No. 09/563,664, filed on May 2, 2000, now Pat. No. 6,505,160.

(60) Provisional application No. 60/178,028, filed on Jan. 26, 2000.

(51) **Int. Cl.⁷** G06K 9/00; H04N 7/167; H04L 9/00

(52) **U.S. Cl.** 382/100; 380/240; 713/168

(58) **Field of Search** 382/100, 232, 382/305, 112, 115, 345/744, 765; 380/54, 201, 210, 240; 386/46, 52, 69; 704/23, 270, 273; 707/104.1, 200; 713/153, 176, 168; 725/37

(56) **References Cited**

U.S. PATENT DOCUMENTS

3,493,674 A	2/1970	Houghton
3,569,619 A	3/1971	Simjian
3,576,369 A	4/1971	Wick et al.
3,585,290 A	6/1971	Sanford
3,655,162 A	4/1972	Yamamoto et al.

3,703,628 A	11/1972	Philipson, Jr.
3,809,806 A	5/1974	Walker et al.
3,838,444 A	9/1974	Loughlin et al.
3,914,877 A	10/1975	Hines
3,922,074 A	11/1975	Ikegami et al.
3,971,917 A	7/1976	Maddox et al.
3,977,785 A	8/1976	Harris
3,982,064 A	9/1976	Barnaby
3,984,624 A	10/1976	Waggener
4,025,851 A	5/1977	Haselwood et al.
4,184,700 A	1/1980	Greenaway
4,225,967 A	9/1980	Miwa et al.
4,230,990 A	10/1980	Lert, Jr. et al.
4,231,113 A	10/1980	Blasbalg
4,238,849 A	12/1980	Gassmann

(List continued on next page.)

FOREIGN PATENT DOCUMENTS

CA	2235002	12/1998
DE	3806411	9/1989
DE	19521969 C1	2/1997

(List continued on next page.)

OTHER PUBLICATIONS

U.S. patent application Ser. No. 09/314,648, Rodriguez et al., filed May 19, 1999.

(List continued on next page.)

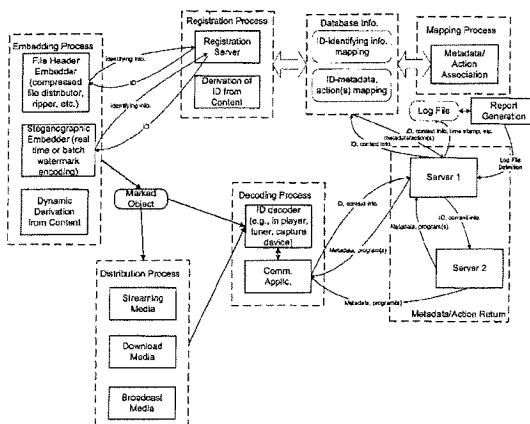
Primary Examiner—Jayanti K. Patel

(74) *Attorney, Agent, or Firm*—Digimare Corporation

(57) **ABSTRACT**

Media objects are transformed into active, connected objects via identifiers embedded into them or their containers. In the context of a user's playback experience, a decoding process extracts the identifier from a media object and possibly additional context information and forwards it to a server. The server, in turn, maps the identifier to an action, such as returning metadata, re-directing the request to one or more other servers, requesting information from another server to identify the media object, etc. The linking process applies to broadcast objects as well as objects transmitted over networks in streaming and compressed file formats.

29 Claims, 3 Drawing Sheets



US 6,829,368 B2

Page 2

U.S. PATENT DOCUMENTS

4,252,995 A	2/1981	Schmidt et al.	4,967,273 A	10/1990	Greenberg
4,262,329 A	4/1981	Bright et al.	4,969,041 A	11/1990	O'Grady et al. 348/473
4,296,326 A	10/1981	Haslop et al. 283/70	4,972,471 A	11/1990	Gross et al.
4,297,729 A	10/1981	Steynor et al. 360/40	4,972,475 A	11/1990	Sant'Anselmo
4,313,197 A	1/1982	Maxemchuk 370/210	4,972,476 A	11/1990	Nathans 713/186
4,367,488 A	1/1983	Leventer et al. 370/204	4,977,594 A	12/1990	Shear 705/53
4,379,947 A	4/1983	Warner 370/204	4,979,210 A	12/1990	Nagata et al.
4,380,027 A	4/1983	Leventer et al. 348/467	4,996,530 A	2/1991	Hilton
4,389,671 A	6/1983	Posner et al.	5,003,590 A	3/1991	Lechner et al.
4,395,600 A	7/1983	Lundy et al. 381/73.1	5,010,405 A	4/1991	Schreiber et al.
4,416,001 A	11/1983	Ackerman	5,023,907 A	6/1991	Johnson 710/200
4,423,415 A	12/1983	Goldman	5,027,401 A	6/1991	Soltesz 380/54
4,425,642 A	1/1984	Moses et al. 370/477	5,034,982 A	7/1991	Heninger et al.
4,476,468 A	10/1984	Goldman	5,036,513 A	7/1991	Greenblatt
4,523,508 A	6/1985	Mayer et al.	5,040,059 A	8/1991	Leberl 348/135
4,528,588 A	7/1985	Löfberg 340/5.1	5,053,956 A	10/1991	Donald 713/601
4,547,804 A	10/1985	Greenberg 348/460	5,062,666 A	11/1991	Mowry et al. 283/67
4,553,261 A	11/1985	Froessl	5,063,446 A	11/1991	Gibson
4,590,366 A	5/1986	Rothfjell	5,073,899 A	12/1991	Collier et al.
4,595,950 A	6/1986	Lofberg	5,073,925 A	12/1991	Nagata et al.
4,618,257 A	10/1986	Bayne et al. 356/71	5,075,773 A	12/1991	Pullen et al.
4,637,051 A	1/1987	Clark	5,077,608 A	12/1991	Dubner
4,639,779 A	1/1987	Greenberg	5,077,795 A	12/1991	Rourke et al.
4,647,974 A	3/1987	Butler et al.	5,079,648 A	1/1992	Maufe
4,654,867 A	3/1987	Labedz et al.	5,091,966 A	2/1992	Bloomberg et al.
4,660,221 A	4/1987	Dlugos	5,095,196 A	3/1992	Miyata 235/382
4,663,518 A	5/1987	Borror et al.	5,103,459 A	4/1992	Gilhousen et al. 370/206
4,665,431 A	5/1987	Cooper	5,113,437 A	5/1992	Best
4,672,605 A	6/1987	Hustig et al. 370/201	5,113,445 A	5/1992	Wang 380/51
4,675,746 A	6/1987	Tetrick et al. 358/296	5,128,525 A	7/1992	Stearns et al.
4,677,435 A	6/1987	Cause D'Agraves et al.	5,144,660 A	9/1992	Rose
4,682,794 A	7/1987	Margolin	5,146,457 A	9/1992	Veldhuis et al. 370/523
4,703,476 A	10/1987	Howard	5,148,498 A	9/1992	Resnikoff et al.
4,712,103 A	12/1987	Gotanda	5,150,409 A	9/1992	Elsner
4,718,106 A	1/1988	Weinblatt	5,161,210 A	11/1992	Druyvesteyn et al.
4,723,149 A	2/1988	Harada	5,166,676 A	11/1992	Milheiser
4,739,377 A	4/1988	Allen 355/133	5,168,146 A	12/1992	Marshall et al.
4,750,173 A	6/1988	Bluthgen 370/528	5,181,786 A	1/1993	Hujink 400/61
4,765,656 A	8/1988	Becker et al.	5,185,736 A	2/1993	Tyrrell et al.
4,775,901 A	10/1988	Nakano	5,199,081 A	3/1993	Saito et al.
4,776,013 A	10/1988	Katri et al.	5,200,822 A	4/1993	Bronfin et al. 725/22
4,805,020 A	2/1989	Greenberg	5,212,551 A	5/1993	Conanan
4,807,031 A	2/1989	Broughton et al. 348/460	5,213,337 A	5/1993	Sherman 463/40
4,811,357 A	3/1989	Betts et al.	5,216,724 A	6/1993	Suzuki et al. 382/135
4,811,408 A	3/1989	Goldman	5,228,056 A	7/1993	Schilling
4,820,912 A	4/1989	Samyn	5,243,411 A	9/1993	Shirochi et al.
4,835,517 A	5/1989	van der Gracht et al.	5,243,423 A	9/1993	DeJean et al. 348/473
4,855,827 A	8/1989	Best 348/485	5,245,165 A	9/1993	Zhang
4,864,618 A	9/1989	Wright et al.	5,245,329 A	9/1993	Gokcebay
4,866,771 A	9/1989	Bain	5,247,364 A	9/1993	Banker et al.
4,874,936 A	10/1989	Chandler et al.	5,253,078 A	10/1993	Balkanski et al.
4,876,617 A	10/1989	Best et al.	5,257,119 A	10/1993	Funada et al.
4,879,747 A	11/1989	Leighton et al. 713/186	5,259,025 A	11/1993	Monroe 705/75
4,884,139 A	11/1989	Pommier	5,267,334 A	11/1993	Normille et al.
4,885,632 A	12/1989	Mabey et al.	5,280,537 A	1/1994	Sugiyama et al. 370/529
4,888,798 A	12/1989	Earnest 705/54	5,288,976 A	2/1994	Citron 235/375
4,903,301 A	2/1990	Kondo et al.	5,291,243 A	3/1994	Heckman et al. 399/3
4,908,836 A	3/1990	Rushforth et al. 375/152	5,293,399 A	3/1994	Hefti
4,908,873 A	3/1990	Philibert et al. 382/100	5,295,203 A	3/1994	Krause et al. 382/248
4,918,484 A	4/1990	Ujiie et al.	5,299,019 A	3/1994	Pack et al.
4,920,503 A	4/1990	Cook	5,305,400 A	4/1994	Butera
4,921,278 A	5/1990	Shiang et al.	5,315,098 A	5/1994	Tow 235/494
4,939,515 A	7/1990	Adelson	5,319,453 A	6/1994	Copriviza et al.
4,941,150 A	7/1990	Iwasaki	5,319,724 A	6/1994	Blonstein et al.
4,943,973 A	7/1990	Werner	5,319,735 A	6/1994	Preuss et al.
4,943,976 A	7/1990	Ishigaki	5,321,470 A	6/1994	Hasuo et al. 399/366
4,944,036 A	7/1990	Hyatt 367/43	5,325,167 A	6/1994	Melen
4,947,028 A	8/1990	Gorog 235/380	5,327,237 A	7/1994	Gerdes et al.
4,963,998 A	10/1990	Maufe	5,337,362 A	8/1994	Gormish et al.
4,965,827 A	10/1990	McDonald	5,349,655 A	9/1994	Mann
			5,351,302 A	9/1994	Leighton et al.

US 6,829,368 B2

Page 3

5,371,792 A	12/1994	Asai et al.	5,629,770 A	S/1997	Brassil		
5,374,976 A	12/1994	Spannenburg	399/366	5/1997	Stefik et al.		
5,379,345 A	1/1995	Greenberg	455/2.01	6/1997	Rhoads	382/232	
5,387,941 A	2/1995	Montgomery et al.		6/1997	Stefik	705/54	
5,394,274 A	2/1995	Kahn		6/1997	Rubin		
5,396,559 A	3/1995	McGrew		6/1997	Wellner	725/100	
5,398,283 A	3/1995	Virga		7/1997	Saito	705/54	
5,404,160 A	4/1995	Schober et al.		7/1997	Kawakami et al.	348/463	
5,404,377 A	4/1995	Moses		8/1997	Schmid	235/375	
5,408,542 A	4/1995	Callahan		8/1997	Kawana		
5,416,307 A	5/1995	Danck et al.	235/449	5,663,766 A	9/1997	Sizer, II	348/473
5,418,853 A	5/1995	Kanota et al.		5,664,018 A	9/1997	Leighton	380/54
5,422,963 A	6/1995	Chen et al.		5,665,951 A	9/1997	Newman et al.	235/375
5,422,995 A	6/1995	Aoki et al.		5,666,487 A	9/1997	Goodman et al.	
5,425,100 A	6/1995	Thomas et al.		5,668,636 A	9/1997	Beach et al.	358/296
5,428,606 A	6/1995	Moskowitz		5,671,282 A	9/1997	Wolff et al.	713/179
5,428,607 A	6/1995	Hiller et al.	370/352	5,673,316 A	9/1997	Auerbach et al.	705/51
5,428,731 A	6/1995	Powers	707/501.1	5,687,236 A	11/1997	Moskowitz et al.	380/28
5,432,542 A	7/1995	Thibadeau et al.		5,710,636 A	1/1998	Curry	358/3.28
5,432,870 A	7/1995	Schwartz		5,719,939 A	2/1998	Tel	713/179
5,446,488 A	8/1995	Vogel		5,721,788 A	2/1998	Powell et al.	382/100
5,450,122 A	9/1995	Keene		5,727,092 A	3/1998	Sandford, II et al.	382/251
5,450,490 A	9/1995	Jensen et al.		5,735,547 A	4/1998	Morelle et al.	283/67
5,461,426 A	10/1995	Limberg et al.		5,740,244 A	4/1998	Indeck et al.	713/176
5,463,209 A	10/1995	Figh	235/383	5,742,845 A	4/1998	Wagner	395/821
5,469,222 A	11/1995	Sprague	348/580	5,745,604 A	4/1998	Rhoads	382/232
5,469,506 A	11/1995	Berson et al.	713/186	5,761,686 A	6/1998	Bloomberg	707/529
5,473,631 A	12/1995	Moses		5,765,152 A	* 6/1998	Erickson	707/9
5,479,168 A	12/1995	Johnson et al.		5,768,426 A	6/1998	Rhoads	382/232
5,481,294 A	1/1996	Thomas et al.		5,774,452 A	6/1998	Wolosewicz	
5,488,664 A	1/1996	Shamir		5,778,102 A	7/1998	Sandford, II et al.	382/251
5,493,677 A	2/1996	Bfalogh	707/104.1	5,790,693 A	8/1998	Graves et al.	382/135
5,495,581 A	2/1996	Tsai	707/526	5,790,697 A	8/1998	Munro et al.	382/135
5,496,071 A	3/1996	Walsh	283/70	5,804,803 A	9/1998	Cragun et al.	235/375
5,499,294 A	3/1996	Friedman		5,809,160 A	9/1998	Powell et al.	382/100
5,502,576 A	3/1996	Ramsay et al.	358/444	5,809,317 A	9/1998	Kogan et al.	707/501.1
5,515,081 A	5/1996	Vasilik		5,817,205 A	10/1998	Kaule	382/294
5,521,722 A	5/1996	Colvill et al.	358/500	5,818,441 A	10/1998	Throckmorton et al.	345/717
5,524,933 A	6/1996	Kunt et al.		5,819,289 A	10/1998	Sanford, II et al.	707/104.1
5,530,751 A	6/1996	Morris		5,825,871 A	10/1998	Mark	379/357.03
5,530,759 A	6/1996	Braudaway et al.	380/54	5,825,892 A	10/1998	Braudaway et al.	380/51
5,530,852 A	6/1996	Meske, Jr. et al.	709/206	5,838,458 A	11/1998	Tsai	358/402
5,532,920 A	7/1996	Hartrick et al.		5,841,978 A	11/1998	Rhoads	709/217
5,537,223 A	7/1996	Curry		5,848,144 A	12/1998	Ahrens	379/219
5,539,471 A	7/1996	Myhrvold et al.		5,848,413 A	12/1998	Wolff	707/10
5,539,735 A	7/1996	Moskowitz		5,852,673 A	12/1998	Young	382/164
5,541,662 A	7/1996	Adams et al.		5,857,038 A	1/1999	Odawa et al.	382/284
5,544,255 A	8/1996	Smithies et al.		5,862,218 A	1/1999	Steinberg	713/176
5,548,646 A	8/1996	Aziz et al.		5,862,260 A	1/1999	Rhoads	382/232
5,557,333 A	9/1996	Jungo et al.		5,869,819 A	2/1999	Knowles et al.	235/375
5,559,559 A	9/1996	Jungo et al.		5,871,615 A	2/1999	Harris	162/140
5,568,179 A	10/1996	Diehl et al.		5,872,589 A	2/1999	Morales	725/24
5,568,550 A	10/1996	Ur	382/306	5,875,249 A	2/1999	Mintzer et al.	380/54
5,568,570 A	10/1996	Rabbani		5,892,900 A	4/1999	Ginter et al.	
5,572,010 A	11/1996	Petrie		5,893,101 A	4/1999	Balogh et al.	707/100
5,572,247 A	11/1996	Montgomery		5,898,779 A	4/1999	Squilla et al.	713/176
5,576,532 A	11/1996	Hecht		5,900,608 A	5/1999	Iida	235/381
5,579,124 A	11/1996	Aijala et al.		5,902,353 A	5/1999	Reber et al.	709/219
5,582,103 A	12/1996	Tanaka et al.		5,903,729 A	5/1999	Reber et al.	709/219
5,587,743 A	12/1996	Montgomery		5,903,892 A	* 5/1999	Hoffert et al.	707/10
5,590,197 A	12/1996	Chen et al.		5,905,248 A	5/1999	Russell et al.	235/462.15
5,594,226 A	1/1997	Steger	235/379	5,905,251 A	5/1999	Knowles	235/472.01
5,598,526 A	1/1997	Daniel et al.	345/540	5,905,810 A	5/1999	Jones et al.	382/135
5,602,920 A	2/1997	Bestler et al.		5,913,210 A	6/1999	Call	707/4
5,606,609 A	2/1997	Houser et al.		5,915,027 A	6/1999	Cox et al.	380/54
5,611,575 A	3/1997	Petrie		5,930,767 A	7/1999	Reber et al.	705/26
5,613,004 A	3/1997	Cooperman et al.	380/28	5,932,863 A	8/1999	Rathus et al.	235/462.15
5,613,012 A	3/1997	Hoffman et al.		5,933,798 A	8/1999	Linnartz	702/91
5,614,940 A	3/1997	Cobbley et al.		5,933,829 A	8/1999	Durst et al.	707/10
5,617,119 A	4/1997	Briggs et al.	707/100	5,938,726 A	8/1999	Reber et al.	709/217
5,617,148 A	4/1997	Montgomery		5,938,727 A	8/1999	Ikeda	709/218

US 6,829,368 B2

Page 4

5,939,695 A	8/1999	Nelson	235/383	GB	2204984	11/1988
5,940,595 A	8/1999	Reber et al.	709/227	JP	4-248771	2/1992
5,949,055 A	9/1999	Fleet et al.	235/469	JP	5/242217	9/1993
5,950,173 A	9/1999	Perkowski	705/26	JP	8-30759	2/1996
5,963,916 A	10/1999	Kaplan	705/26	WO	WO 89/08915	9/1989
5,971,277 A	10/1999	Cragun et al.	235/462.01	WO	WO 93/25038	12/1993
5,974,141 A	10/1999	Saito	705/52	WO	WO94/27228	11/1994
5,974,548 A	10/1999	Adams	713/200	WO	WO95/04665	2/1995
5,978,773 A	11/1999	Hudetz et al.	705/23	WO	WO95/10813	4/1995
5,979,757 A	11/1999	Tracy et al.	235/383	WO	WO 95/10835	4/1995
5,983,176 A	* 11/1999	Hoffert et al.	704/233	WO	WO 95/14289	5/1995
5,983,218 A	11/1999	Syeda-Mahmoo	707/3	WO	WO 95/20291	7/1995
5,991,426 A	11/1999	Cox et al.	382/100	WO	WO 96/26494	8/1996
6,005,501 A	12/1999	Wolosewicz	341/52	WO	WO 96/27259	9/1996
6,024,287 A	2/2000	Takai et al.		WO	WO97/43736	11/1997
6,035,177 A	3/2000	Moses et al.	725/22	WO	WO98/14887	4/1998
6,052,486 A	4/2000	Knowlton et al.	382/232	WO	WO98/20642	5/1998
6,064,764 A	5/2000	Bhaskaran et al.	382/183	WO	WO98/24050	6/1998
6,122,403 A	9/2000	Rhoads	382/233	WO	WO98/40823	9/1998
6,166,750 A	12/2000	Negishi	347/131	WO	WO98/49813	11/1998
6,188,787 B1	2/2001	Ohmiae et al.		WO	WO99/34277	7/1999
6,243,480 B1	6/2001	Zhao et al.		WO	WO99/36876	7/1999
6,266,430 B1	7/2001	Rhoads et al.	382/100	WO	WO00/44131	7/2000
6,286,036 B1	9/2001	Rhoads		WO	WO 01/08405	2/2001
6,301,360 B1	10/2001	Bocioneck et al.		WO	WO 01/08169	10/2001
6,311,214 B1	10/2001	Rhoads 382/100
6,321,648 B1	11/2001	Benson et al.				
6,321,981 B1	11/2001	Ray et al.				
6,324,573 B1	11/2001	Rhoads				
6,324,574 B1	11/2001	Gong	709/218			
6,343,204 B1	1/2002	Yang				
6,359,985 B1	3/2002	Koch et al.				
6,411,725 B1	6/2002	Rhoads				
6,496,802 B1	12/2002	van Zoest et al.				
6,522,769 B1	2/2003	Rhoads et al.				
6,542,927 B2	4/2003	Rhoads				
6,553,129 B1	4/2003	Rhoads				
6,577,746 B1	6/2003	Evans et al.				
2001/0017709 A1	8/2001	Murakami et al.				
2001/0024510 A1	9/2001	Iwamura				
2001/0026618 A1	10/2001	Van Wie et al.				
2001/0026629 A1	10/2001	Oki				
2001/0030759 A1	10/2001	Hayashi et al.				
2001/0053299 A1	12/2001	Matsunoshita et al.				
2002/0001095 A1	1/2002	Kawakami et al.				
2002/0003891 A1	1/2002	Hoshino				
2002/0009208 A1	1/2002	Alattar et al.	382/100			
2002/0018228 A1	2/2002	Torigoe				
2002/0051237 A1	5/2002	Ohara				

FOREIGN PATENT DOCUMENTS

EP	366381 A2	10/1989
EP	372 601	6/1990
EP	411 232	2/1991
EP	418 964 A1	3/1991
EP	441 702	8/1991
EP	493 091	7/1992
EP	058 482	8/1992
EP	551 016	7/1993
EP	581 317	2/1994
EP	605 208	7/1994
EP	649 074	4/1995
EP	705 025	4/1996
EP	711061	5/1996
EP	0789480	8/1997
EP	872995	10/1998
EP	0642060 B1	4/1999
EP	1122939	8/2001
GB	2063018	5/1981
GB	2067871	7/1981
GB	2196167	4/1988

U.S. patent application Ser. No. 09/342,688, Rodriguez et al., filed Jun. 29, 1999.

U.S. patent application Ser. No. 09/342,971, Rodriguez et al., filed Jun. 29, 1999.

U.S. patent application Ser. No. 09/679,261, Davis et al., filed Oct. 4, 2000.

U.S. patent application Ser. No. 09/562,517, Davis et al., filed May 1, 2000.

U.S. patent application Ser. No. 09/547,664, Rhoads et al., filed Apr. 12, 2000.

U.S. patent application Ser. No. 09/571,442, Rhoads et al., filed May 15, 2000.

U.S. patent application Ser. No. 09/858,189, Rhoads et al., filed May 14, 2001.

U.S. patent application Ser. No. 09/631,409, Brundage et al., filed Aug. 3, 2000.

U.S. patent application Ser. No. 09/452,021, Davis et al., filed Nov. 30, 1999.

U.S. patent application Ser. No. 09/629,401, Seder et al., filed Aug. 1, 2000.

U.S. patent application Ser. No. 09/473,396, Evans et al., filed Dec. 28, 1999.

U.S. patent application Ser. No. 09/563,664, Levy et al., filed May 2, 2000.

U.S. patent application Ser. No. 09/670,115, Rhoads et al., filed Sep. 26, 2000.

Brassil et al., Electronic Marking and Identification Techniques to Discourage Document Copying, Proceedings of INFOCOM '94 Conference on Computer, IEEE Commun. Soc Conference, Jun. 12-16, 1994, 1278-1287.

Bruckstein, A.M.; Richardson, T.J., A holographic transform domain image watermarking method, Circuits, Systems, and Signal Processing vol. 17, No.3 p. 361-89, 1998. This paper includes an appendix containing an internal memo of Bell Labs, which according to the authors of the paper, was dated Sep. 1994.

"High Water FBI Limited Presentation Image Copyright Protection Software," FBI Ltd brochure, Jul., 1995, 17 pp.

Koch et al., "Copyright Protection for Multimedia Data," Fraunhofer Institute for Computer Graphics, Dec. 16, 1994, 15 pp.

Koch et al., "Towards Robust and Hidden Image Copyright Labeling," Proc. of 1995 IEEE Workshop on Nonlinear Signal and Image Processing, Jun. 20-22, 1995, 4 pp.

Kurak et al., "A Cautionary Note On Image Downgrading," 1992 IEEE, pp. 153-159.

Mintzer et al., "Safeguarding Digital library Contents and Users" Digital Watermarking, D-Lib Magazine, Dec. 1997; ISSN 1082-9873.

Rindfrey, "Towards an Equitable System for Access Control and Copyright Protection in Broadcast Image Services: The Ecuicrypt Approach," Intellectual Property Rights and New Technologies, Proc. of the Conference, R. Oldenbourg Verlag Wien Munchen 1995, 12 pp.

Schreiber et al., "A Compatible High-Definition Television System Using the Noise-Margin Method of Hiding Enhancement Information," SMPTE Journal, Dec. 1989, pp. 873-879.

SDMI Example Use Scenarios (Non-Exhaustive), Version 1.2, Jun. 16, 1999.

Szepanski, "A Signal Theoretic Method for Creating Forgery-Proof Documents for Automatic Verification," Proceedings 1979 Carnahan Conference on Crime Countermeasures, May 16, 1979, pp. 101-109.

Szepanski, "Additive Binary Data Transmission for Video Signals," Papers Presented at Conf. Of Comm. Engineering Soc. Sep. 30-Oct. 3, 1980, Technical Reports vol. 74, pp. 342-352.

Tanaka et al., "A Visual Retrieval System with Private Information for Image Database," Proceeding International Conference on DSP Applications and Technology, Oct. 1991, pp. 415-421.

Tanaka et al., "New Integrated Coding Schemes for Computer-Aided Facsimile," Proc. IEEE Int'l Conf. on Sys. Integration, Apr. 1990, pp. 275-281.

Tirkel et al., "Electronic Water Mark," DICTA-93, Macquarie University, Syndney, Australia, Dec., 1993, pp. 666-673.

Weber et al., "Correlative Image Registration," Seminars in Nuclear Medicine, vol XXIV, No. 4, Oct., 1994, pp. 311-323.

Szepanski, "A Signal Theoretic Method for Creating Forgery-Proof Documents for Automatic Verification," Proceedings 1979 Carnahan Conference on Crime Countermeasures, May 16, 1979, pp. 101-109.

Dautzenberg, "Watermarking Images," Department of Microelectronics and Electrical Engineering, Trinity College Dublin, 47 pp., Oct. 1994.

Szepanski, "Additive Binary Data Transmission for Video Signals," Conference of the Communications Engineering Society, 1980, NTG Technical Reports, vol. 74, pp. 343-351. (German text and English translation enclosed).

U.S. patent application Ser. No. 09/404,291, Levy, filed Sep. 23, 1999.

U.S. patent application Ser. No. 09/234,780, Rhoads/Gustafson, filed Jan. 20, 1999.

U.S. patent application Ser. No. 09/478,713, Cookson, filed Jan. 6, 2000.

Cookson, Chris, General Principles of Music Uses on Portable Devices, presented to SDMI, Mar. 5, 1999.

Winograd, J.M., "Audio Watermarking Architecture for Secure Digital Music Distribution," a Proposal to the SDMI Portable Devices Working Group, by Aris Technologies, Inc., Mar. 26, 1999.

Mintzer et al., "Safeguarding Digital Library Contents and Users: Digital Watermarking," D-Lib Magazine, Dec. 1997, 12 pp.

U.S. patent application Ser. No. 09/765,102, Shaw, filed Jan. 17, 2001.

U.S. patent application Ser. No. 09/761,349, Rhoads, filed Jan. 16, 2001.

U.S. patent application Ser. No. 09/761,280, Rhoads, filed Jan. 16, 2001.

U.S. patent application Ser. No. 09/645,779, Tian et al., filed Aug. 24, 2000.

U.S. patent application Ser. No. 09/689,226, Brunk, filed Oct. 11, 2000.

U.S. patent application Ser. No. 09/689,250, Ahmed, filed Oct. 11, 2000.

U.S. patent application Ser. No. 09/689,293, Tian et al., filed Oct. 11, 2000.

U.S. patent application Ser. No. 09/625,577, Carr et al., filed Jul. 25, 2000.

U.S. patent application Ser. No. 09/574,726, Rhoads et al., filed May 18, 2000.

U.S. patent application Ser. No. 09/562,524, Carr et al., May 1, 2000.

U.S. patent application Ser. No. 09/498,223, Rhoads et al., filed Feb. 3, 2000.

U.S. patent application Ser. No. 09/465,418, Rhoads et al., filed Dec. 16, 1999.

U.S. patent application Ser. No. 09/431,990, Rhoads, filed Nov. 3, 1999.

U.S. patent application Ser. No. 09/428,359, Davis et al., filed Oct. 28, 2000.

U.S. patent application Ser. No. 09/342,972, Rhoads, filed Jun. 29, 1999.

U.S. patent application Ser. No. 09/293,602, Rhoads, filed Apr. 15, 1999.

U.S. patent application Ser. No. 09/293,601, Rhoads, filed Apr. 15, 1999.

U.S. patent application Ser. No. 09/287,940, Rhoads, filed Apr. 7, 1999.

U.S. patent application Ser. No. 09/185,380, Davis et al., filed Nov. 3, 1998.

U.S. patent application Ser. No. 09/074,034, Rhoads, filed May 6, 1998.

U.S. patent application Ser. No. 09/127,502, Rhoads, filed Jul. 31, 1998.

Audio Watermarking Architectures for Secure Digital Music Distribution, A Proposal to the SDMI Portable Devices Working Group by ARIS Technologies, Inc, Mar. 26, 1999, pp. 1-11.

Audio Watermarking Architectures for Persistent Protection, Presentation to SDMI PDWG, Mar. 29, 1999, J. Winegrad, Aris Technologies, pp 1-16.

Audio Watermarking System to Screen Digital Audio Content for LCM Acceptance, A Proposal Submitted in Response to PDWG99050504-Transition Cfp by ARIS Technologies, Inc., May 23, 1999, Document Version 1.0, 15 pp.

Boland et al., "Watermarking Digital Images for Copyright Protection", *Fifth Int'l Conference on Image Processing and it's Application*, Jul. 1995, pp. 326-330.

Levy, "AIPL's Proposal for SDMI. An Underlying Security System" (slide presentation), Mar. 29, 1999, 23 slides.

Microsoft Response to CfP for Technology Solutions to Screen Digital Audio Content for LCM Acceptance, SDMI, PDWG Tokyo, May 23, 1999, 9 pp.

Response to CfP for Technology Solutions to Screen Digital Audio Content for LCM Acceptance, NTT Waveless Radio Consortium, May 23, 1999, 9 pp.

Sandford II et al., "The Data Embedding Method", *Proceedings of the SPIE* vol. 2615, pp. 226-259, 1996.

Thomas, Keith, *Screening Technology for Content from Compact Discs*, May 24, 1999, 11 pp.

Tirkel et al., "Electronic Water Mark," *Dicta-93*, Marquarie University, Sydney, Australia, Dec., 1993, pp. 666-672.

Vidal et al., "Non-Noticeable Information Embedding in Color Images: Marking and Detection", *IEEE* 1999, pp. 293-297.

Wolfgang et al., "A Watermark for Digital Images," *Computer Vision and Image Processing Laboratory, Purdue University*, Sep. 1996, pp. 219-222.

"Access Control and COpright Protection for Images, WorkPackage 8: Watermarking," Jun. 30, 1995, 46 pp.

"Access Control and COpright Protection for Images, WorkPackage 3: Evaluation of Existing Systems," Apr. 19, 1995, 68 pp.

"Access Control and COpright Protection for Images, WorkPackage 1: Access Control and Copyright Protection for Images Need Evaluation," Jun., 1995, 21 pp.

"Access Control and COpright Protection for Images, Conditional Access and Copyright Protection Based on the Use of Trusted Third Parties," 1995, 43 pp.

Arachelian, "White Noise Storm," Apr. 11, 1994, Internet reference, 13 pp.

Arazi, et al., "Intuition, Perception, and Secure Communication," *IEEE Transactionson Systems, Man, and Cybernetics*, vol. 19, No. 5, Sep./Oct. 1989, pp. 1016-1020.

Arthur, "Digital Fingerprints Protect Artwork," *New Scientist*, Nov. 12, 1994, p. 24.

Aura, "Invisible Communication," Helsinki University of Technology, Digital Systems Laboratory, Nov. 5, 1995, 13 pp.

Bender et al., "Techniques for Data Hiding," Draft Preprint, Private Correspondence, dated Oct. 30, 1995.

Bender et al., "Techniques for Data Hiding," Massachusetts Institute of Technology, Media Laboratory, Jan. 1995, 10 pp.

Boneh, "Collusion-Secure Fingerprinting for Digital Data," Department of Computer Science, Princeton University, 1995, 31 pp.

Boney et al., "Digital Watermarks for Audio Signals," *Proceedings of Multimedia '96*, 1996 IEEE, PP. 473-480.

Boucqueau et al., Equitable Conditional Access and Copyright Protection for Image Based on Trusted Third Parties, Teleservices & Multimedia Communications, 2nd Int. Cost 237 Workshop, Second International Cost 237 Workshop, Nov., 1995; published 1996, pp. 229-243.

Brassil et al., "Hiding Information in Document Images," Nov., 1995, 7 pp.

Brown, "S-Tools for Windows, Version 1.00, .COPYRGT. 1994 Andy Brown, What is Steganography," Internet reference, Mar. 6, 1994, 6 pp.

Bruyndonckx et al., Neural Network Post-Processing of Coded Images Using Perceptual Masking, 1994, 3 pp.

Bruyndonckx et al., "Spatial Method for Copyright Labeling of Digital Images," 1994, 6 pp.

Burgett et al., "A Novel Method for Copyright Labeling Digitized Image Data," requested by e-mail from author (unavailable/password protected on IGD WWW site); received Sep. 18, 1995, 12 pp.

Caronni, "Assuring Ownership Rights for Digital Images," Published in the Proceedings of 'Reliable IT Systems,' VIS '95, HH. Bruggemann and W. Gerhardt-Hackl (Ed.), Vieweg Publishing Company, Germany, 1995, Jun. 14, 1994, 10 pp.

Caruso, "Digital Commerce, 2 plans for watermarks, which can bind proof of authorship to electronic works." *New York Times*, Aug. 7, 1995, one p.

Castro et al., "Registration of Translated and Rotated Images Using Finite Fourier Transforms," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. PAMI-9, No. 5, Sep. 1987, pp. 700-703.

Choudhury, et al., "Copyright Protection for Electronic Publishing over Computer Networks," *IEEE Network Magazine*, Jun. 1994, 18 pp.

Clarke, "Invisible Code Tags Electronic Images," *Electronic Engineering Times*, Jun. 12, 1995, n. 852, p. 42.

"Copyright Protection for Digital Images, Digital Fingerprinting from FBI," Highwater FBI brochure, 1995, 4 pp.

"The Copyright Can of Worms Opened Up By The New Electronic Media," *Computergram Internations*, pCGN07170006, Jul. 17, 1995 and "The Copyright Can of Worms Opened Up By The New Electronic Media—2," *Computergram Internations*, pCGN07210008, Jul. 21, 1995, 3 pp. total.

Cox et al., "Secure Spread Spectrum Watermarking for Multimedia," *NEC Research Institute Technical Report*, Dec. 5, 1995, 33 pp.

Cox et al., "A Secure, Imperceptible Yet Perceptually Salient, Spread Spectrum Watermark for Multimedia," *IEEE, Southcon'96, Conference Record*, Pp. 192-197, 1996.

"Cyphertech Systems: Introduces Digital Encoding Device to Prevent TV Piracy," *Hollywood Reporter*, Oct. 20, 1993, p. 23.

Delaigle et al., "Digital Watermarking," *Proc. SPIE—Int. Soc. Opt. Eng.*, vol. 2659, pp. 99-110, 1996.

Delaigle et al., "A Psychovisual Approach for Digital Picture Watermarking," 1995, 20 pp.

DICE Digital Watermark System, Q&A, Dec., 1995, 12 pp.

Digimare presentation at RSA Conference, approximately Jan. 17, 1996, 4 pp.

Fimmerstad, "The Virtual Art Museum," *Ericsson Connexion*, Dec., 1995, pp. 29-31.

Fitzgerald, "Invisible Digital Copyright ID," *Editor & Publisher*, Jun. 25, 1994, p. 62.

"Foiling Card Forgers With Magnetic 'Noise,'" *Wall Street Journal*, Feb. 8, 1994.

Frequently Asked Questions About Digimare Signature Technology, Aug. 1, 1995, <HTTP://WWW.DIGIMARC.COM>, 9 pp.

Friedman, "The Trustworthy Digital Camera: Restoring Credibility to the Photographic Image," *IEEE Transactions on Consumer Electronics*, vol. 39, No. 4, Nov., 1993, pp. 905-910.

Gabor, et al., "Theory of Communication," *J. Inst. Elect. Eng.* 93, 1946, pp. 429-441.

Hartung et al., Digital Watermarking of Raw and Compressed Video, *Proc. SPIE 2952, Digital Compression Technologies and Systems for Video Communications*, Oct., 1996, pp 205-213.

Hecht, "Embedded Data Glyph Technology for Hardcopy Digital Documents," SPIE vol. 2171, Feb. 1994, pp. 341–352.

"Holographic signatures for digital images," The Seybold Report on Desktop Publishing, Aug. 1995, one p.

Humphrey, "Stamping Out Crime," Hollywood Reporter, Jan. 26, 1994, p. S48.

Jain, "Image Coding Via a Nearest Neighbors Image Model," IEEE Transactions on Communications, vol. COM-23, No. 3, Mar. 1975, pp. 318–331.

Johnson, "Steganography," Dec. 10, 1995, 32 pp.

JPEG Group's JPEG Software (release 4), FTP.CSU-A.BERKELEY.EDU/PUB/CYPHERPUNKS/APPLICATIONS/JSTEG/JPEG.ANNOUNCEMENT.GZ.

Kassam, Signal Detection in Non-Gaussian Noise, Dowden & Culver, 1988, pp. 1–96.

Koch et al., "Digital Copyright Labeling: Providing Evidence of Misuse and Tracking Unauthorized Distribution of Copyrighted Materials," Oasis Magazine, Dec. 1995, 3 pp.

Luc, "Analysis of Spread Spectrum System Parameters for Design of Hidden Transmission," Radioengineering, vol. 4, No. 2, Jun. 1995, pp. 26–29.

Machado, "Announcing Stego 1.0a2, The First Steganography Tool for the Macintosh," Internet reference, Nov. 28, 1993, 3 pp.

Macq, "Cryptography for Digital TV Broadcasting," Proceedings of the IEEE, vol. 83, No. 6, Jun. 1995, pp. 944–957.

Matthews, "When Seeing is Not Believing," New Scientist, Oct. 16, 1993, pp. 13–15.

Matsui et al., "Video-Steganography: How to Secretly Embed a Signature in a Picture," IMA Intellectual Property Project Proceedings, Jan. 1994, vol. 1, Issue 1, pp. 187–205.

Mintzer et al., "Toward on-line, Worldwide Access to Vatican Library Materials," IBM J. Res. Develop. vol. 40 No. 2, Mar., 1996, pp. 139–162.

Moller, et al., "Rechnergestutzte Steganographie: Wie sie Funktioniert und warum folglich jede Reglementierung von Verschlüsselung unsinnig ist," DuD, Datenschutz und Datensicherung, Jun. 18, 1994 318–326.

"NAB—Cyphertech Starts Anti-Piracy Broadcast Tests," Newsbytes, NEW03230023, Mar. 23, 1994.

Nakamura et al., "A Unified Coding Method of Image and Text Data Using Discrete Orthogonal Transform," Systems and Computers in Japan, vol. 21, No. 3, 1990, pp. 87–92.

Nakamura et al., "A Unified Coding Method of Dithered Image and Text Data Using Micropatterns," Electronics and Communications in Japan, Part 1, vol. 72, No. 4, 1989, pp. 50–56.

New Product Information, "FBI at AppleExpo" (Olympia, London), Nov., 1995, 2 pp.

Ohnishi et al., Embedding a Seal into a Picture Under Orthogonal Wavelet Transform, Proceedings of Multimedia '96, 1996, IEEE, pp. 514–421.

ORuanaidh et al., "Watermarking Digital Images for Copyright Protection," <http://www.kalman.mee.tcd.ie/people/jjr/eva.sub.—pap.html>, Feb. 2, 1996, 8 pp. (Also published Aug., 1996, IEEE Proceedings—Vision, Image and Signal Processing, vol. 143, No. 4, pp. 250–256.).

Pennebaker et al., JPEG Still Image Data Compression Standard, Chapter 3, "Aspects of the Human Visual System," pp. 23–27, 1993, Van Nostrand Reinhold, New York.

Pickholtz et al., "Theory of Spread-Spectrum Communications—A Tutorial," Transactions on Communications, vol. COM-30, No. 5, May, 1982, pp. 855–884.

Pitas et al., "Applying Signatures on Digital Images," IEEE Workshop on Nonlinear Image and Signal Processing, Neos Marmaras, Greece, pp. 460–463, Jun., 1995.

Port, "Halting Highway Robbery on the Internet," Business Week, Oct. 17, 1994, p. 212.

Roberts, "Picture Coding Using Pseudorandom Noise," IRE Trans. on Information Theory, vol. 8, No. 2, Feb., 1962, pp. 145–154.

Sapwater et al., "Electronic Copyright Protection," Photo>Electronic Imaging, vol. 37, No. 6, 1994, pp. 16–21.

Schneier, "Digital Signatures, Cryptographic Algorithms Can Create Nonforgeable Signatures for Electronic Documents, Making Them Valid Legal Instruments" BYTE, Nov. 1993, pp. 309–312.

shaggy@phantom.com, "Hide and Seek v. 4.0," Internet reference, Apr. 10, 1994, 3 pp.

Short, "Steps Toward Unmasking Secure Communications," International Journal of Bifurcation and Chaos, vol. 4, No. 4, 1994, pp. 959–977.

Simmons, "Subliminal Channels; Past and Present," ETT, vol. 5, No. 4, Jul.–Aug. 1994, pp. 45–59.

Sheng et al., "Experiments on Pattern Recognition Using Invariant Fourier–Mellin Descriptors," Journal of Optical Society of America, vol. 3, No. 6, Jun., 1986, pp. 771–776.

Sklar, "A Structured Overview of Digital Communications—a Tutorial Review—Part I," IEEE Communications Magazine, Aug., 1983, pp. 1–17.

Sklar, "A Structured Overview of Digital Communications—a Tutorial Review—Part II," IEEE Communications Magazine, Oct., 1983, pp. 6–21.

"Steganography," Intellectual Property and the National Information Infrastructure The Report of the Working Group on Intellectual Property Rights, Sep. 1995, pp. 212–213.

Tanaka et al., "Embedding Secret Information Into a Dithered Multi-Level Image," Proc. IEEE Military Comm. Conf., Sep. 1990, pp. 216–220.

Tanaka, "Embedding the Attribute Information Into a Dithered Image," Systems and Computers in Japan, vol. 21, No. 7, 1990, pp. 43–50.

Tirkel et al., "A Two-Dimensional Digital Watermark," 1995, 6 pp.

Toga et al., "Registration Revisited," Journal of Neuroscience Methods, 48 (1993), pp. 1–13.

van Schyndel et al., "Towards a Robust Digital Watermark," ACCV '95, vol. 2, Dec., 1995, pp. 504–508.

Wagner, "Fingerprinting," 1983 IEEE, pp. 18–22.

Walton, "Image Authentication for a Slippery New Age," Dr. Dobb's Journal, Apr. 1995, pp. 18–26, 82–87.

"Watermarking & Digital Signature: Protect Your Work!" Published on Internet 1996, <http://Itswww.epfl.ch/about.jordan/watermarking.html>.

Wise, "The History of Copyright, Photographers' Rights Span Three Centuries," Photo>Electronic Imaging, vol. 37, No. 6, 1994.

van Schyndel et al., "A Digital Watermark," IEEE International Conference on Image Processing, Nov. 13–16, 1994, pp. 86–90.

Zhao et al., "Embedding Robust Labels Into Images for Copyright Protection," Proc. of the International Congress on Intellectual Property Rights for Specialized Information, Knowledge and New Technologies (Vienna, Austria) Aug. 21–25, 1995, 10 pp.

Bender, "Applications for Data Hiding," IBM Systems Journal, vol. 39, No. 3–4, pp. 547–568, 2000.

US 6,829,368 B2

Page 8

Gruhl et al., "Information Hiding to Foil the Casual Counterfeiter," Proc. 2d Information Hiding Workshop, LNCS vol. 1525, pp. 1-15 (Apr. 15, 1998).
U.S. patent application Ser. No. 60/000,442, Hudetz, filed Jun. 20, 1995.
U.S. patent application Ser. No. 60/082,228, Rhoads, filed Apr. 16, 1998.
U.S. patent application Ser. No. 60/141,763, Davis, filed Jun. 30, 1999.
U.S. patent application Ser. No. 60/158,015, Davis et al., filed Oct. 6, 1999.
U.S. patent application Ser. No. 60/071,983, Levy, filed Jan. 20, 1998.

U.S. patent application Ser. No. 09/404,291, Levy, filed Sep. 23, 1999.

U.S. patent application Ser. No. 60/114,725, Levy, filed Dec. 31, 1998.

U.S. patent application Ser. No. 09/234,780, Rhoads/Gustafson, filed Jan. 20, 1999.

U.S. patent application Ser. No. 60/116,641, Cookson, filed Jan. 21, 1999.

U.S. patent application Ser. No. 09/478,713, Cookson, filed Jan. 6, 2000.

* cited by examiner

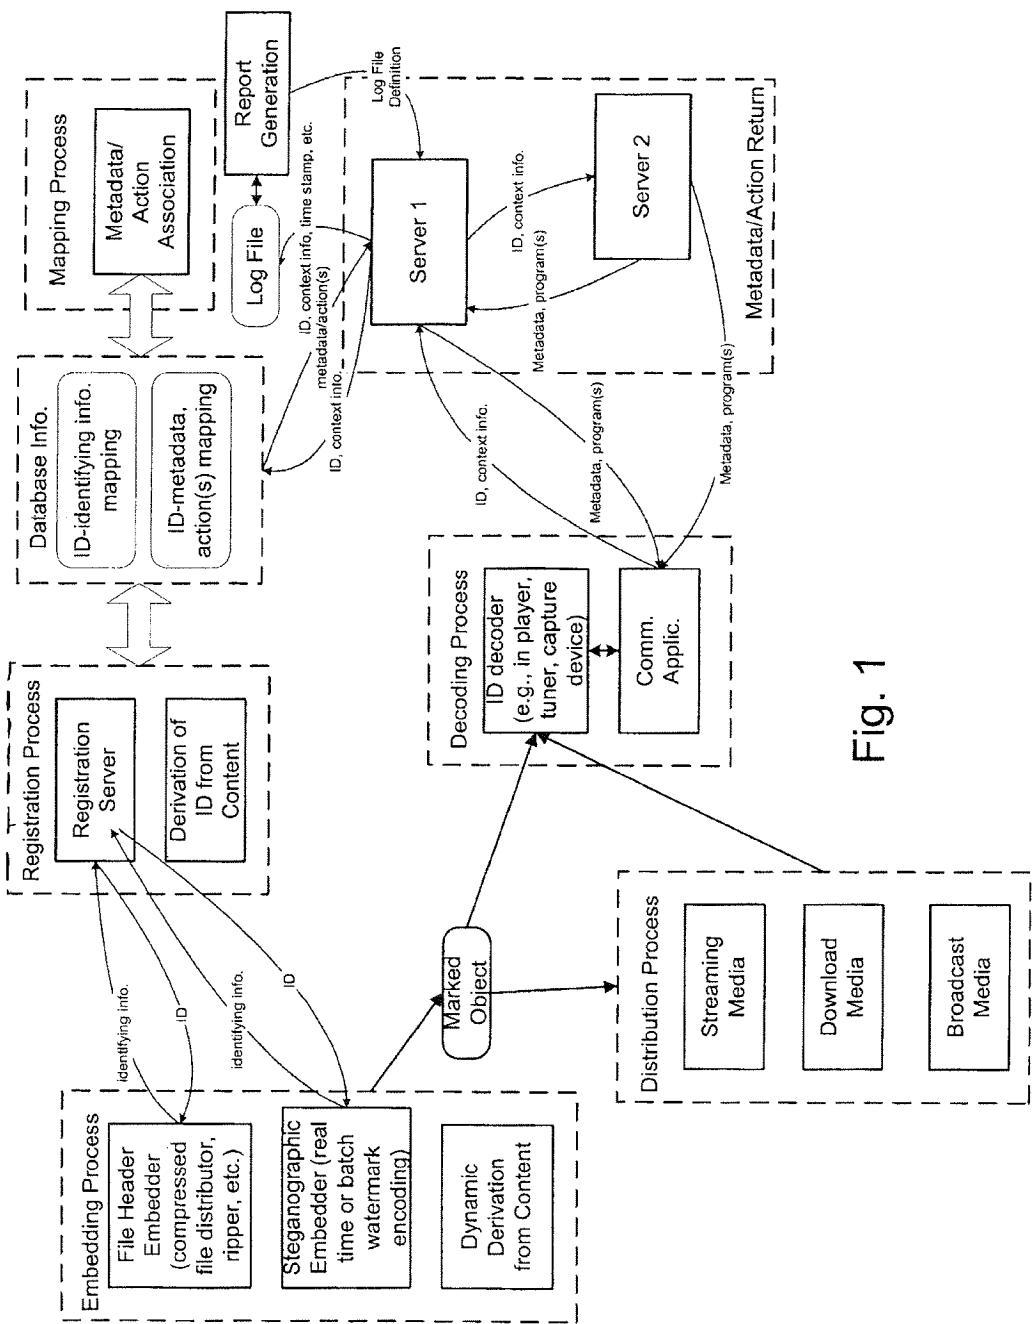


Fig. 1

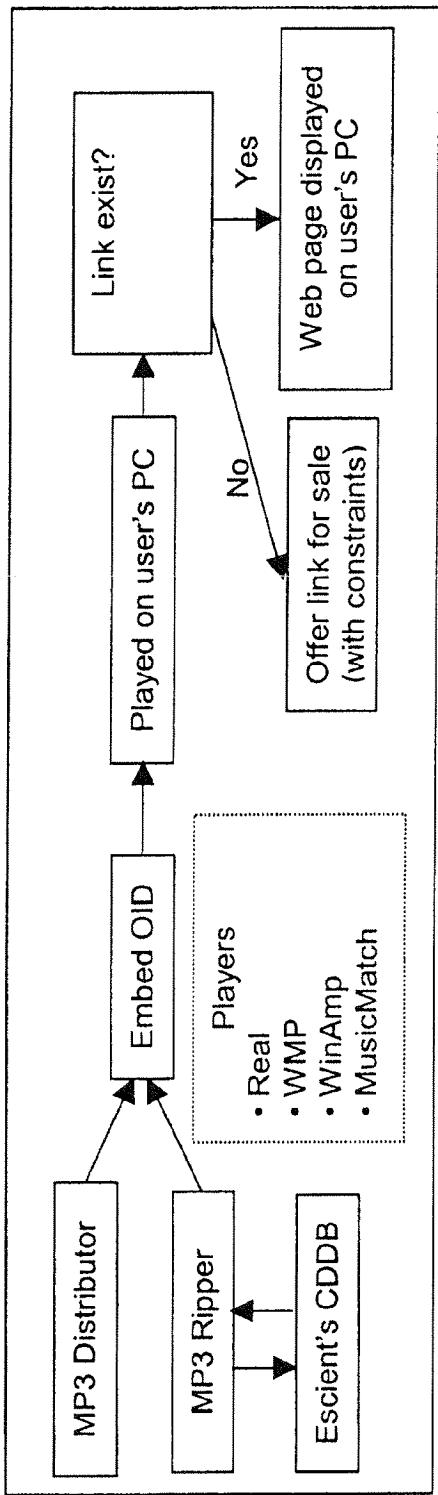
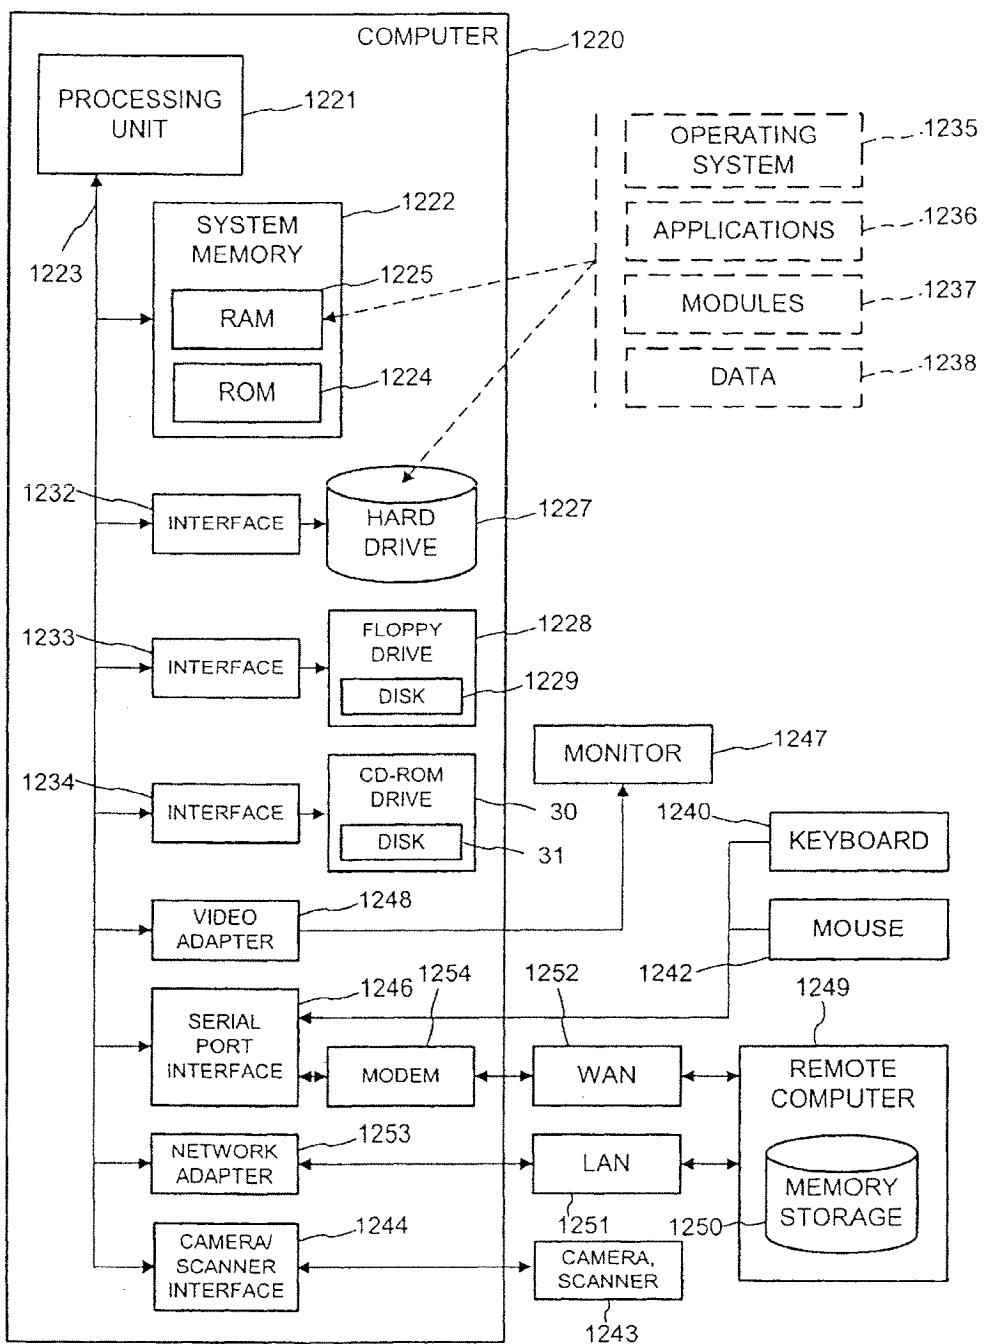


Fig. 2

Fig. 3



ESTABLISHING AND INTERACTING WITH ON-LINE MEDIA COLLECTIONS USING IDENTIFIERS IN MEDIA SIGNALS

RELATED APPLICATION DATA

This patent application claims the benefit of U.S. Provisional Application No. 60/178,028, filed Jan. 26, 2000. This patent application is also a continuation-in-part of U.S. patent application Ser. No. 09/563,664, filed May 2, 2000 (Now U.S. Pat. No. 6,505,160). These applications are hereby incorporated by reference.

The subject matter of the present application is related to that disclosed in U.S. Pat. Nos. 5,862,260 and 6,122,403; PCT Applications Nos. WO 01/01331 A1, published Jan. 4, 2001, and WO 00/70585, published Nov. 23, 2000; and in co-pending U.S. Pat. application Ser. Nos. 09/292,569, filed Apr. 15, 1999; Ser. No. 09/343,104, filed Jun. 29, 1999; Ser. No. 09/473,396, filed Dec. 28, 1999 (Now U.S. Pat. No. 6,577,746); Ser. No. 09/476,686, filed Dec. 30, 1999; Ser. No. 09/503,881, filed Feb. 14, 2000 (Now U.S. Pat. No. 6,614,914); Ser. No. 09/525,865, filed Mar. 15, 2000 (Now U.S. Pat. No. 6,611,607); Ser. No. 09/547,664, filed Apr. 12, 2000; Ser. No. 09/574,726, filed May 18, 2000; and Provisional Application No. 60/191,778, filed Mar. 24, 2000. Each of these documents is hereby incorporated by reference.

TECHNICAL FIELD

The invention relates to linking audio and other multimedia data objects with metadata and actions via a communication network, e.g., computer, broadcast, wireless, etc.

BACKGROUND AND SUMMARY

Developments in network technology and media content (e.g., images, video, and audio) storage, delivery, and playback are re-shaping the entertainment, information technology, and consumer electronics industries. With these developments, there are an increasing number of applications for associating media content with auxiliary data. The auxiliary data may provide information describing the content, copy control information or instructions, links to related content, machine instructions, etc. This auxiliary data is sometimes referred to as metadata. In many applications, metadata suffers from the drawback that it is vulnerable to becoming separated from an associated media signal.

Steganography provides a way to embed data in the media signal. As such, it offers an advantage over conventional ways to associate metadata with media signals. Examples of steganography include digital watermarking and data glyphs. Exemplary watermarking techniques suitable for still image and video content are shown in U.S. Pat. No. 5,862,260 to Rhoads and U.S. Pat. No. 5,915,027 to Cox. Exemplary watermarking techniques suitable for use with audio content are shown in the just-cited Rhoads patent, as well as U.S. Pat. No. 5,945,932 to Smith and U.S. Pat. No. 5,940,135 to Petrovic.

Advances in computer and wireless networking, multimedia coding, and higher bandwidth communication links are creating many new ways to distribute and enjoy multimedia content, such as music and movies. Coding formats for audio like MPEG 1 Layer 3 (MP3) have already caused significant changes in music delivery to consumers. Despite the advances in technology, content distributors and broadcasters still need to address how to effectively promote and sell content.

This document describes systems and processes for linking audio and other multimedia data objects with metadata and actions via a communication network, e.g., computer, broadcast, wireless, etc. Media objects are transformed into active, connected objects via identifiers embedded into them or their containers. These identifiers can be embedded by the owner or distributor of the media object, or automatically created from the media object. In the context of a user's playback experience, a decoding process extracts the identifier from a media object and possibly additional context information and forwards it to a server. The server, in turn, maps the identifier to an action, such as returning metadata, re-directing the request to one or more other servers, requesting information from another server to identify the media object, etc. If the identifier has no defined action, the server can respond with an option for the user to buy the link and control the resulting action for all objects with the current identifier. The linking process applies to broadcast objects as well as objects transmitted over networks in streaming and compressed file formats.

Further features will become apparent with reference to the following detailed description and accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram illustrating examples of media object linking processes and systems.

FIG. 2 is a diagram illustrating media object linking applications.

FIG. 3 is a diagram illustrating an operating environment for multimedia content management and delivery applications.

DETAILED DESCRIPTION

Linking Audio and other Media Objects via Identifiers

The following sections describe systems and processes for linking audio and other media objects to metadata and actions via an identifier. For the sake of illustration, the disclosure focuses on a specific media type, namely audio signals (e.g., music, sound tracks of audio visual works, voice recordings, etc.). However, these systems, their components and processes apply to other types of media signals as well, including video, still images, graphical models, etc. As described further below, an identifier attached to an audio signal is used to connect that signal with metadata and/or programmatic or device actions. In the context of this document, the terms "media object" and "audio object" refer to an electronic form of a media signal and audio signal, respectively. The linking of media signals applies to objects that are transmitted over wire networks (such as a computer network), wireless networks (such as a wireless telephone network), and broadcast (AM, FM, digital broadcast, etc.).

There are a number of ways to associate an identifier with an audio object. One way to associate the identifier is to insert it in the form of a numeric or alphanumeric code (e.g., binary or M-ary code) in the electronic file in which the audio is stored. Another way to associate the identifier is to embed it as auxiliary data in the audio signal using steganographic methods, such as digital watermarking or other data hiding techniques. Yet another way is to derive the identifier from the audio signal, the table of contents, the file system structure, or its container (e.g., an electronic file or physical package for data like flash memory, Digital Versatile Disk (DVD), minidisk, or compact disk (CD). The physical media may have identifying characteristics, such as a unique identifier or encoded metadata, or other attributes from which an identifier can be derived (e.g., CD disk wobble).

When the identifier is associated with metadata or actions, it transforms the media object into a "linked" object. The identifier travels with the object through distribution, including in some cases, through physical distribution in packaged media and through electronic distribution (broadcast or network communication). The identifier may travel within the same band as the audio object, such as a watermark, or via a separate band, such as a file header or footer or separate broadcast band. A decoding device or programmatic process extracts the identifier from the object and uses it to retrieve related data or actions ("metadata"). In the case of an audio object, like a song, the metadata typically includes the title, artist, lyrics, copyright owner, sound recording owner, information about buying or sampling opportunities and URLs to this type of data as well as web sites and other programs and devices. Linked actions include device or programmatic processes for electronically establishing a license, transferring content (either streaming or download), sending an email, recording marketing data about a transaction, etc. The identifier allows a fan of a particular type of music or artist to get more information about the music and to buy more music. From the perspective of the artists and record labels, the identifier provides an additional opportunity to promote their music and sell content, concert tickets, etc.

In addition, in some implementations where identifier linking transactions are monitored, it enables the vendors of music to gather data about electronic transactions triggered by the link. For example, users of information may choose to provide information about themselves when they register their decoding device or software with the system. A user ID or other context information may then be recorded when the identifier is extracted and used to trigger a transaction. Many entities involved in the distribution of media signals can benefit from the linking capability. Artists can link their music to information about themselves and provide electronic buying opportunities for music, concert tickets, clothing, etc. Rights holding organizations can use the link to inform users about itself and licensing opportunities. In some cases, the link may also be used to monitor playing and distribution of copies of the music. Record labels can link their music to information about the artist, the label, electronic buying opportunities, etc. Electronic retailers can increase sales by linking users to opportunities to sample and buy additional music (via download or streaming delivery over a wire or wireless network). Conventional brick and mortar retailers can use linking to provide information about the music and to provide buying opportunities. Radio stations and other broadcasters can use the linking capability to bring users to their web sites, creating advertising revenue, to provide electronic buying opportunities for music, concert tickets, clothing items, etc. These and other forms of linked metadata and actions may be implemented in various combinations in different application scenarios.

Depending on the application, the identifier may identify the media object in which it is embedded, or entities, things or actions other than that particular media object. One type of identifier is an object ID that identifies an audio object. This identifier may be a number associated with the object, such as its International Standard Recording Code (ISRC). Another type of identifier is distributor ID that identifies the distributor of the audio object. Another type of identifier is a broadcaster ID that identifies the broadcaster of the audio object. Of course, more than one identifier may be encoded into an audio object or its container. In the event that an object ID is not encoded with an audio object, but instead, a distributor or broadcaster identifier is encoded with the object, other context information, such as the time of play

back or distribution, location of distribution, etc. may be used to identify the audio object as part of the linking process. An example is a radio station that marks its broadcasts with a station ID and maintains a playlist database with the air times of each audio object. At decoding time, the station ID is extracted and used along with context information such as the air time of the audio object to look up the audio object or its corresponding metadata and actions. This approach enables the linking system to provide audio object specific metadata or actions even without requiring a unique object identifier in every audio object.

System Implementation

FIG. 1 is a diagram of a system configuration of linked media objects. In this configuration, an identifier links audio objects to metadata via an electronic network, such as the Internet, a wireless network, or a broadcast network. As depicted in FIG. 1, an embedding process may be used to encode an identifier in an audio object or its container. In some cases, an embedding process encodes the identifier in the audio file (e.g., a tag in a file header or footer), in the audio signal (a digital watermark), or in the physical packaging. The identifier may also be derived as a function of the audio signal or other information in the file or physical packaging (e.g., track information on a CD). In the case of dynamically derived identifiers, an embedding process is not necessary because the identifier can be derived from the content at decoding time.

In some application scenarios, the embedding process interacts with a registration process to get an identifier. The embedding process provides information about the object (e.g., a title and artist name, an ISRC, name of distributor, etc.). In response, the registration process provides an identifier and stores a database record of the association between identifier and the object or other information used in decoding to identify the object, such as its distributor or broadcaster. The registration process may be used to assign an identifier to an audio object and to distributors or broadcasters of audio objects. The embedding and registration processes may occur before the audio object is distributed to consumers, or sometime thereafter, such as when a user transfers (e.g., "rips") a media object from one format to another (e.g., a packaged format to an electronic file format such as a compressed file format).

Once registered, an interactive or automated mapping process associates the identifier with data or actions. The registration process creates a database of identifiers and associates the identifiers with corresponding media objects, distributors, broadcasters, etc. The mapping process associates the identifiers with corresponding metadata or actions.

Once associated with an audio object and metadata, the identifier transforms the audio object into a linked object. The identifier remains with the object through distribution, although some embedding processes are more robust than others to intentional or unintentional distortion/removal of the identifier. There are a variety of different distribution scenarios. Some examples depicted in FIG. 1 include transferring an audio object over a computer network, streaming the object over a computer network, or broadcasting it (e.g., AM/FM broadcasting, digital broadcasting, broadcasting over wireless carriers, etc.). Whatever the distribution process, a user ultimately receives the linked object in a player, tuner, or capture device.

To activate the linked object, a decoding process extracts the identifier and uses it to access associated data or actions. The decoding process may be implemented as a separate program or device, or integrated into a player, tuner, or some other capture device, such as a listening devices that con-

verts ambient audio waves to an electronic signal and then extracts the identifier from the signal.

In the configuration shown in FIG. 1, the decoding process forwards the extracted identifier to a communication application, which in turn, forwards it in a message to a server. The decoding process or the communication application may add additional context information to the message sent to the to a server. The context information may relate to the user, the user's device, the attributes of the session (time of playback, format of playback, type of distribution (e.g., broadcast or transmitted audio file), etc.) Based on identifier and optional context information, the server determines an associated action to perform, such as re-directing an identifier or context data to another server, returning metadata (including programs, content, etc.), downloading content, logging a transaction record. To find the associated action or actions, the server maps the identifier to actions based on the information established in the mapping process. The server may: 1) look up the data and actions in a local database stored in its memory subsystem; 2) route the identifier to one or more other servers via the network, which in turn look up related actions and data associated with the identifier; or 3) perform some combination of actions 1 and 2.

In the first case, server 1 returns data or actions associated with the identifier. The server may look up related data based on the identifier alone, or based on the identifier and other context information. Context information may be information provided by the user, by the user's computer or device, or by some other process or device. In the second case, the server looks up one or more addresses associated with the identifier and forwards the identifier and/or possibly other context data to secondary servers at these addresses via conventional networking protocols. Again, this context data may include data from the user, the user's computer, some other device or database. For example, server 1 might query a remote database for instructions about how to process an identifier. These instructions may specify data to return to the communication application or to forward to another server, which in turn, looks up associated data and returns it to the communication application. A server may return data that an audio player displays to the user or uses to control rendering of the content. For example, the server can tell the player that the object contains inappropriate content for children. The player or user can make decisions about whether or how to play the material based on this information.

Both the server and the player can adopt a set of rules. The server rules may be used to control what the server returns in response to an identifier and context data. The player rules may be used to control what the player displays to the user or how it renders the content based on data returned from a server.

Either the first server, or a server one or more levels of indirection from the identifier may return data and programmatic actions to a player via the communication application. Each server in these levels of indirection receives a database key, such as an identifier or context information, from the previous server, and uses it to look up corresponding actions. These actions may include returning data or programs to the communication application or to previous servers in the routing path of the message from the communication application. Also, the servers may route requests for information or actions to other servers. The server or servers may return data or perform actions in response to the identifier (or other context data) that do not directly impact the decoding process, or the device in which it operates.

The system depicted in FIG. 1 allows several different interested parties to establish services linked via the identifier. For example, server 1 can be configured to provide generic promotional and/or licensing information associated with an identifier. If the content owner, distributor, retailer, artist or other related party wishes to provide information or services for a connected object, then server 1 may also route the identifier for that object, and possibly context information, the address of the communication application, and instructions, to servers maintained by these entities. These servers, in turn, provide promotional, sales, or licensing information, and electronic buying or licensing opportunities specific to that entity back to the consumer over the network via the communication application.

In the context of a network configuration, Internet protocols may be used to return data to the communication application or to the device or system in which it operates. The communication application may be implemented in a web browser, such as Internet Explorer or Netscape Navigator. Examples of ways of exchanging information between a client player and a server include returning a web page with metadata and program scripts designed to run on the end user's system. The metadata itself may include active links, such as URLs to other network resources, such as a web site or some other network service. The path of the identifier from the decoding process, and the return path from a server to the communication application may include one or more hops through a wire or wireless connection using standard wire and wireless communication protocols like TCP/IP, HTTP, XML, WAP, Bluetooth, etc. In addition, data returned to the user may be routed through one or more servers that may forward the data, and in some cases, augment the data or modify it in some fashion.

FIG. 2 is a diagram illustrating applications of the system depicted in FIG. 1. In the application scenarios depicted in FIG. 2, an embedding process encodes an object identifier (OID) into an audio file, such as an ID3 tag in the header of an MP3 file or audio frame headers in the MP3 file. FIG. 2 shows two embedding scenarios. The first is an MP3 distributor that embeds OIDs in MP3 files before transmitting them over a network, such as the Internet, typically via a web site interface. The second is a file ripping process where a programmed computer or other device extracts an audio object from packaged media such as a CD and converts it into a coded file format like MP3. In the latter case, the ripping process may extract metadata from the CD, such as the table of contents, and use this metadata as a key to a database (CDDB) to get information about the songs on the CD, such as title, artists, etc. The table of contents or other metadata from a package medium, such as optical or magnetic storage or flash memory, may be hashed into an index to a database entry that stores information about the media signal stored on the medium. The ripping process uses the information returned from the database to identify the audio objects on the packaged media so that they can be associated with an OID. This is an example of identifying information used to associate an OID with an audio object. As part of the coding process, the ripping process inserts the OID in the file header of the MP3 file.

Later, when a user opens or plays the marked MP3 in a player, such as a software player like the real player, Liquid Audio player, Windows Media Player (WMP), WinAmp, MusicMatch, etc., a plug-in software module in the player extracts the OID and forwards it to a server via an Internet connection. The plug-in may establish its own Internet connection, or pass the OID to an Internet Browser, which in turn, establishes a connection (if one is not already

present) with the server. As an intermediate step, the plug-in may display a window with user options, such as "learn more about the song", "play the song", or both. The user can then choose to get more information by actuating the first or third options in the user interface window, which cause the plug-in to forward the OID to the server.

The server then returns a web page associated with the OID, or re-directs the OID to another server (e.g., one maintained by the content distributor or owner), which in turn, returns a web page of information about the object and links to related actions (e.g., a link to a licensing server, a link to a server for buying and downloading related music etc.). The licensing server may be programmed to download software players and new music offerings compatible with those players. For instance, the licensing server may provide software for decrypting, decoding, and playing electronically distributed music according to usage rules packaged with the electronically distributed music. In this application scenario, the linking of the MP3 file enables the content owner to market music and products that promote the sale of audio objects in other formats, included formats protected with encryption, watermark copy managements schemes, etc.

In the event that a media object is not linked, the decoding and server processes can be programmed to enable the user to purchase a link for the object. For example in one scenario, the player plug-in displays a graphic for a link information indicating that the link is available after determining that an OID is not in the file. If the user clicks on the graphic, the plug-in displays more information about the procedure for purchasing or renting a link. This information may be provided in conjunction with querying the server and displaying information returned from the server, or alternatively, providing pre-programmed information incorporated into the plug-in. If the user is interested in purchasing the link, he or she can then enter input (e.g., click on a button such as "Get Link") that initiates the process of registering an OID with the object and associating metadata or actions with the OID. The process of registering the OID and associating the OID with metadata or actions may be performed as described in this document. This scenario provides yet another mechanism for transforming content into connected content.

There are many possible variations to the applications scenarios illustrated in FIG. 2. During the file ripping process (or some other embedding process), the embedder may generate a unique ID from the metadata read from the packaged media on which the media object resides. One example of such an ID is the number derived from CD metadata currently used to index information in the Cddb database. This ID may then be embedded in the audio object or its file header/footer. During OID registration, the registration process may inform the embedding process that the OID (and thus, the object for which it was derived) has not been associated with metadata or actions. In this case, the user may be given an opportunity to purchase the link, either at the time of ripping, or in the future, wherever the object travels. In the latter case, the OID in the object is associated with an option to buy the link and customize the data and/or actions associated with that link. Rather than link to promotional information, the OID gives users an option to buy or rent the link and provides them with an opportunity to customize it (e.g., linking it to a custom web site). Once customized, other users that open or play the file will then be able to link to the customized information or actions.

To assert control over the type of customization that users may perform, the registration and mapping processes can

place constraints on the types of metadata and actions that users can link to a media object.

In the multimedia content industry, there are typically many rights holders and entities involved in the distribution process. This may present a conflict when linking a media object to one entity. One way to address this problem is have an object link to many different entities. For example, the server could map an OID to many entities and return links to retailers, distributors, record labels and artists. Another way to address it is to encode additional information about the distributor in the OID. For example, the OID includes fields that identify the object and its distributor. If a user activates the link to purchase products, including media objects, then the distributor name is logged with the purchase and that distributor is credited with royalties associated with the transaction. The distributor field may also be used as a key to look up the appropriate action for the OID, such as re-directing the OID to the web server of the entity associated with that OID. In this approach, even if the OID directs a user to a record label's website, the distributor field can be used to credit the distributor with a royalty for the linking transaction.

The entity responsible for maintaining a web site linked via on identifier can make deals with online resources for providing data about a media object such as lyrics, song titles, radio station play lists. The website may link to this information, access it via a database manager, etc.

File Identifiers

One form of identifier is an identifier that is inserted in an audio object file, but in a distinct field from the audio signal itself. Some examples are file headers and footers. This file identifier may be assigned before or after distribution of the audio object to consumers. In addition, it may be derived from the audio signal or other information in the file. For example, an identifier generator may derive a unique or sufficiently unique identifier from a portion of a music signal. A variety of methods for generating a unique numbers based on a unique collection of numbers may be used.

The process of embedding a file identifier may be done at the time of encoding or transcoding a file. For example, the file identifier may be inserted during a ripping process, such as when a device or programmatic process converts a song from a format stored on packaged media, like a CD or DVD, to an electronic, and compressed form, such as MP3 or some other audio codec. As another example, the file identifier may be inserted when a device or programmatic process transcodes an electronic music file from one codec format to another. Yet another example is where a file is taken from a digital or analog uncompressed format, and placed in another format for distribution.

Identifiers Embedded in Audio Signal

Another way to associate an identifier with an audio signal is to embed the identifier in the audio signal using steganographic methods, such as digital watermarking or other data hiding techniques. Many of such techniques have been developed and are described in published articles and patents. Watermarking methods are described in U.S. Pat. No. 6,614,914. Other examples of methods for encoding and decoding auxiliary signals into audio signals include U.S. Pat. Nos. 5,862,260, 5,940,135 and 5,945,932. For more information on steganographic applications, see the patent applications and patents incorporated by reference.

The steganographic embedding method may be performed in a batch process. Consider a distributor of electronic music via the Internet or some other network, or a broadcaster of music such as a radio station. In each case, the distributor and broadcaster have a collection of audio

objects. The embedding process may operate on this collection of objects in a batch process by retrieving an electronic version, encoding an identifier obtained from the registration process, and returning the marked version for later distribution or broadcasting. In some cases, it is desirable to do watermark embedding in an iterative process in a studio environment to encode the watermark with an intensity that achieves desired perceptibility and robustness requirements.

The steganographic embedding method may also be performed at the time of transmission of an electronic file or broadcast of the audio object. In the case of distribution via a network such as the Internet (e.g., streaming or file download), real time embedding enables the embedding process to also embed context information that is specific to the consumer (or the consumer's computer) that has electronically ordered the object. For example, when the user requests a file in a streaming or a compressed file format via the Internet using her browser, the distributor's server can request information (perhaps voluntary) about the user to be associated with the transmitted object. Later, the decoding process or the servers that map the identifier to actions or metadata can use this information to determine the types of information to provide or responsive action to perform.

In the case of broadcasting, real time embedding enables the identifier to be steganographically embedded throughout an electronic version of the audio signal just before, or as part of the broadcasting process.

An object or distributor ID (as well as other identifiers or context information) can be embedded in the payload of a watermark that is also used for copy control. Portion of the watermark can be used to control whether the object can be played, transferred, recorded, etc., while another part can be used to carry identifiers and other metadata for linking functions described in this document. Alternatively, entirely separate watermark encoding and decoding methods may be used for copy control and linking functions.

A watermarking process may be used to encode different watermarks in the various channels of an audio signal. Message information may be embedded in one or more channels, while synchronization or orientation signals used to detect and decode the message information may be encoded in other channels. Also, different messages (e.g., different identifiers) may be encoded in different channels. At decoding time, the different identifiers can trigger different actions or link to different data.

In broadcasting applications, an identifier may be encoded along with the broadcast of the associated media signal by modulating a subcarrier of the main carrier frequency used to transmit the media signal. The subcarrier conveys auxiliary data such as the identifier, while the main carrier conveys the associated media signal. To reduce audibility of the auxiliary data (e.g., the identifier(s)) encoded in the sub-carrier, the data can be randomized by applying it to a pseudorandom or random number by some function that may be inverted in the decoding process, e.g., multiplication or exclusive OR functions. One example of sub-carrier encoding and decoding is Active HSDS 97 developed by Seiko Corporation.

Identifiers in Digital Radio Broadcasts

Some forms of digital radio broadcasts support transmission of metadata along with media signals. This metadata can also be used to carry one or more identifiers that are mapped to metadata or actions. The metadata can be encoded at the time of broadcast or prior to broadcasting. Decoding of the identifier may be performed at the digital receiver. In particular, the digital receiver receives the broadcast data, extracts the identifier, and either automatically, or

at the user's direction, forwards the identifier to a server to look up the associated metadata or action.

Dynamic Identifier Extraction from Audio Content or Related Data

As noted above, another way to associate an identifier with a corresponding audio signal is to derive the identifier from the signal. This approach has the advantage that the embedding process is unnecessary. Instead, the decoding process can generate the identifier from the audio object. In

this case, the decoder computes a fingerprint of the audio signal based on a specified fingerprinting algorithm. The fingerprint is a number derived from a digital audio signal that serves as a statistically unique identifier of that signal, meaning that there is a high probability that the fingerprint was derived from the audio signal in question. One component of fingerprint algorithm is a hash algorithm. The hash algorithm may be applied to a selected portion of a music file (e.g., the first 10 seconds) to create a fingerprint. It may be applied to discrete samples in this portion, or to attributes

that are less sensitive to typical audio processing. Examples of less sensitive attributes include most significant bits of audio samples or a low pass filtered version of the portion. Examples of hashing algorithms include MD5, MD2, SHA, and SHA1.

As an aside, fingerprinting may also be used to determine whether an audio signal has been watermarked. The fingerprinting application can evaluate a fingerprint for a received object and compare it with one for a watermarked object (or unmarked object) to determine whether the object is likely to be watermarked. Certain fingerprints can be associated with certain types of watermark methods. Using the fingerprint, a decoding device can select an appropriate watermark decoding system for the object.

While specifically discussed in the context of audio objects, the fingerprinting process applies to other types of multimedia content as well, including still images, video, graphics models, etc. For still images and video, the identifier can be derived dynamically from a compressed or uncompressed version of the image or video signal. The fingerprinting process may be tuned to generate a specific identifier based on the type of file format. For example, the process extracts the file format from the file (e.g., from a header or footer), then uses a fingerprinting process tailored for that type of file (e.g., a hash of a compressed image or video frame). The dynamic identifier computed by this process may be associated with metadata and/or actions using the processes and systems described in this document.

Registration Process

One way to implement the registration process is to build client and server application programs that communicate over a computer network using standard network communication protocols. The client may be implemented as a software program that provides identifying information about an audio object. It can obtain the information by prompting the user for the identifying information, or from extracting it from the audio object or its container. The server may be implemented as a database management program that manages identifiers and corresponding audio objects. When queried to provide an identifier for particular identifying information, the program checks whether it has already assigned an identifier to an object based on the identifying information. If so, it returns that identifier that has already been assigned. If not, it assigns a new identifier number, creates a new entry in the database for that number and its associated identifying information.

The type of identifier used to link audio objects varies with the application. As such, the registration process may

vary as well. One type of identifier is a unique identifier for an audio object. Another type of identifier is one that identifies some attribute of the audio object, but does not uniquely identify it, such as a distributor or broadcaster identifier. This type of identifier requires additional context information to uniquely identify the audio object at the time of linking it to actions or metadata. For these types of identifiers, the registration process provides information identifying the attribute of the audio object, such as its distributor or broadcaster. In response, the server provides an identifier that may be embedded in several audio objects that share that attribute.

One example is a broadcaster ID, such as a radio station ID. Audio broadcast by the radio station is embedded with this radio station ID. To identify the object, context information such as the play time captured at the tuner is used along with the radio station ID extracted from the received audio signal to identify the audio object. The decoding process forwards this information to a server. Using the radio station ID and context information, the server maps the ID to an appropriate action. This may include querying a radio station's playlist database for an object identifier based on the station ID and context information. The server can then map the object identifier to an action or metadata based on the object ID returned from the playlist database. Other scenarios are possible. For example, the server could forward the station ID, context data and decoder address to a radio station server, which in turn, looks up the appropriate action or metadata (e.g., web page) and sends it to the device that decoded the station ID.

Broadcast content can also be associated with object identifiers. One way to implement the identifier assignment process is to allocate a unique set of identifiers with each broadcaster/distributor. Those broadcasters or distributors are then free to assign the identifiers to media objects as they wish. Once they complete the identifier assignment process, they may then associate the identifiers with the metadata or actions in a mapping process.

Embedding Process

The embedding process may be integrated into a software program along with the client of the registration process described in the previous section. This integration of registration and embedding functions is particularly suited to a batch embedder, where processing time required to request an identifier is less of a concern.

In real time embedding, the identifier or identifiers are preferably available for associated audio objects before embedding begins. For example, the identifiers can be maintained in a local database on the embedding computer or device and indexed by object title. Distributor and broadcast identifiers are more straightforward because they may be applied to several different audio objects.

The embedding process may also be implemented in an embedding clearinghouse system. The embedding clearinghouse is a computer or other electronic system that analyzes media objects and embeds one or more links in the media objects. The clearinghouse may be implemented in a server on a network, such as the Internet and operate on content in a "push," "pull," or some combination of push and pull models. In the push model, users and other systems send media objects to the embedding clearinghouse for analysis and embedding. The pull model, the clearinghouse has the capability to search for and gather media objects for embedding and analysis. One example of this pull model is an Internet search process called a spider that crawls the Internet, searching for media objects to analyze and embed with one or more identifying links.

The embedding clearinghouse analyzes a media object (perhaps based on out of band data like a file header or footer) and inserts an identifier. This identifier may link to a metadata and actions, such as re-direction to a web site offering products, services, and information related to the content. The embedding clearinghouse may incorporate search engine technology to execute a key word search based on information from the media object and then associate the media object with a series of related URLs returned from the Internet search. The process may be automatic, or with some user input to select which sub-set of links should be inserted.

The embedding clearinghouse may also offer an identifier embedding services for those wanting to link their media objects with metadata, actions, etc. In this application scenario, the embedding clearinghouse may be implemented as an Internet server that is accessible via a web page using conventional network communication and web protocols. To access the server, users visit a web page using an Internet browser. In exchange for a fee, which may be tendered electronically over the Internet from the user's computer to the server, the server provides an embedding service to embed an identifier into a media object uploaded from the user via the user's computer and Internet connection. The user can select the information to associate with a media object, such as generic identifying information (e.g., title, author, owner), generic licensing information, or special information or actions. The provider of the embedding clearinghouse server hosts the generic information, while the special purpose information and actions are accessed through re-direction. In particular, the provider of the clearinghouse server links the embedded identifier to an address or set of addresses of servers that provide the special information or actions. Then at decoding time, the decoding process sends the identifier to the provider's server, which in turn, redirects the identifier to a secondary server or servers that provide special purpose information or actions (e.g., redirect to a web page of the content owner, download related content, provide electronic licensing services, etc.). Decoding the ID and Embedded Context Data

The implementation details of the decoding process depend on how the identifier is encoded into an audio object or its container. In the case where the identifier is encoded in a file header or footer, the decoder may be a software program or digital hardware that parses the header/footer and forwards it to the communication application. One way to implement this type of decoder is to integrate it into a media player as a plug in program. Examples of media players include Windows Media Player from Microsoft, Liquid Audio player from Liquid Audio, Winamp, Real Player from Real Networks. Preferably, the plug-in gives the user visual feedback that the identifier has been detected and displays a window with options to access more information or actions available via the link. For example, the user can be presented with a user interfaces prompting the user to click for more information or buying opportunities. If the user selects these options, the plug-in forwards the user selections and identifier to the communication application, which forwards them to the server (e.g., server 1, FIG. 1).

In the case where the identifier is steganographically encoded in the audio object, a corresponding decoder extracts the identifier. This type of decoder may be implemented as a plug in to a software player as described in the previous paragraph. It may also be implemented in a tuner for broadcast content, or in a listening device that captures audio from the ambient environment.

In the case where the identifier is derived from the content or container metadata, the decoder captures the pertinent

portion of the audio object, and generates the identifier as described above. This type of decoder can be implemented in a software or hardware player, a tuner, etc.

The decoder may collect identifiers in response to a user request while objects containing these identifiers are being played. For example, when the user is playing music, he may like a song and want to buy it or get more information. This feature may be implemented by building an interface that has a button or voice recognition that enables the user to request information or a buy/license opportunity. Once captured, identifiers can be forwarded along with user instructions to the appropriate server.

However, one particularly useful feature is to enable the user to fetch information and make orders from music as the music is playing. The system described previously supports this feature because the decoding process can forward the identifier or identifiers, embedded context information, or additional context information (user information, play time, broadcast type, file type, player type, operating system type) to the communication application as the music is playing. The user can trigger the linking action by pressing a "fetch" button, or saying fetch to a voice activated input device that causes the decoding device to package a message and invoke the communication application (e.g., Internet browser). In turn, the communication application forwards the message to a server that parses the message and determines the associated action.

The activation of the "fetch it" feature may be made on a handheld device that communicates with a decoding device in a tuner via a wireless connection. For example, a user may press a button on a remote control device, like a key chain, which sends a wireless signal to a receiver in the tuner. The receiver invokes the decoding process. The tuner may also send metadata from the server to the remote control device for display using a similar wireless connection. Infrared or RF transceivers, for example, may be used to communicate the data back and forth.

The decoding device may also provide continuous decoding of identifiers. When the user requests a "fetch," the identifier and context information for the current song may be forwarded to the server. Also, the decoding device may automatically fetch generic information such as song title and artist so that this information is immediately available to the user.

Another possible implementation is to temporarily buffer identifiers extracted from some predetermined number of the most recent songs, titles, etc. These identifiers can be stored along with other metadata, such as a time stamp, to inform the user when they were captured. The user can then select one or more of the items to send to the server for more information or related actions.

These features may be implemented in one or more devices. While the example above discusses a remote control device and a separate tuner with a decoder, these functions may be integrated into a single device, such as a car stereo, phone handset, personal digital assistant, and a variety of other types of players or tuners.

The identifier enables dynamic linking. Dynamic linking enables the identifier encoded with a media object to remain fixed, while the metadata or actions associated with that identifier can be changed. To change the associated metadata, the mapping process edits the identifier database to associate new metadata or actions with an identifier. The mapping process can be automated to change metadata or actions associated with an identifier at periodic intervals or in response to system events. In addition, a user may change the associated metadata or actions interactively at any time.

To facilitate access to the database, a web-based interface can be added to the database.

Dynamically linked data returned from a server to a player environment can be displayed to the user in a variety of ways. One way is to display it in a web page or user interface window of a player. The data can be animated by scrolling it across the visual display. The data can also be displayed in the form of HTML links, which, when activated, cause the download of other data or initiate actions, such as playing streaming content from a server.

Server Types

As discussed elsewhere, the servers used to link identifiers to actions may be programmed to provide a variety of actions including:

- 15 returning data and HTML links (e.g., in the form of an HTML document, scripts, etc.)
- downloading media signals in streaming or file format
- performing an electronic transaction (selling products like CDs, DVDs, concert tickets, etc. via computer transaction using credit cards, digital money, etc.)
- establishing a license to use a linked media object
- re-directing to another server
- 20 performing database look up operations for related information, links, actions
- 25 performing database look up to uniquely identify a media object based on distributor/broadcaster ID and other context information
- creating a transaction log

This is by no means an exhaustive list. Another type of server action is to initiate a process of searching a database, a collection of databases or the Internet for additional information related to a linked media object. This type of search service may be performed continuously and the results associated with the identifier. Then, in response to a request from a decoding process, the server can return a digest of the results with links to web pages for additional information.

Communication Application

The implementation details of the communication application are highly dependent on the type of communication link and protocols used to connect the decoding process to a server. Above, an Internet browser is provided as an example. A browser may be implemented in conventional PCs, handheld devices, wireless phones, stereo systems, set top boxes, etc. However, the communication application need not be based on computer network protocols. For wireless devices, where the marked content is played on wireless carrier frequencies, the communication application can employ wireless communication technology to forward identifiers and context information to servers that map this information to actions or metadata and return it via a wireless carrier frequency to user's handset.

Tracking Transactions and Report Generation

As depicted in FIG. 1 and described above, the servers for mapping identifiers to actions may be programmed to dispense a transaction log into a log file. A report generation process can then enable users to define and request queries of data from the log file based on a particular identifier, a particular type of context information (time frame, geographic location, user demographics, etc.), a particular action, etc.

Capture Devices

As noted above, the decoding process may be implemented in a variety of devices or software that process media objects. These devices and software include programmable devices such as personal computers, personal digital

assistants, telephone handsets, set-top boxes, personal stereos, hi-fi components, tuners, receivers, televisions, etc. as well as hardwired devices that may be incorporated into these systems and devices.

In some contexts, it is useful to implement a recording function. This is particularly true in devices that receive a broadcast or stream of media content and need to capture at least a portion of it to decode an identifier. Examples of these devices are radio receivers, and wireless telephone handsets. The record function may be automatic or user activated. In the latter case, the user actuates an input device to control the record process and optionally the record duration. For example, the user may hear a song that she likes and press record. The device, in turn, records at least a part of the object that is currently being received (an audio, visual or audio visual signal). The user can then decide contemporaneously or at a later time to execute the identifier decoding process on the recorded signal. The recording function can be designed to execute for a pre-determined or user specified duration.

In the case of radio and television tuners/receivers, the record function can be used to capture a media signal as it is received. In the case of a telephone handset, the record function can be used for a variety of functions, such as recording part of a telephone conversation, recording speech or other ambient audio through a microphone, or recording a media signal received by the handset via a wireless communication channel. The recordings can be compressed and stored in local memory on the device. In addition, they may be annotated with metadata about the media signal, such as a time stamp to show time of capture, a location stamp to show location of capture, metadata extracted from the object (in band or out of band data), etc. A global positioning device may provide the location stamp. Some wireless phone systems are capable of computing location of a telephone handset via triangulation. This location data may be used to provide geographic location coordinates or the name of nearby landmark, city name, etc.

The metadata may be displayed on a display device to help the user remember the context of a particular recording. In addition, it may be provided as context information along with an identifier to a server that links the identifier and context information to metadata or actions.

Transmarking

In some applications, it may be useful to convert auxiliary information embedded in a media signal from one format to another. This converting process is referred to as transmarking. Transmarking may include converting an out of band identifier like a tag in a header/footer to a watermark or vice versa. It may also involve converting a message in one watermark format to another. The process involves a decoding operating on an input media object, and an encoding of the decoded information into the media object. It may also involve a process for removing the mark originally in the input object to avoid interference with the newly inserted mark.

There are a variety of reasons to perform transmarking. One is to make the embedded information more robust to the types of processing that the media object is likely to encounter, such as converting from one watermark used in packaged media to another watermark used in compressed, and electronically distributed media, or a watermark used in radio or wireless phone broadcast transmission applications.

This type of transmarking process may be performed at various stages of a media object's distribution path. As suggest previously, an identifier in a watermark or file header/footer may be encoded at the time of packaging the

content for distribution, either in an electronic distribution format or a physical packaged medium, such as an optical disk or magnetic memory device. At some point, the media signal may be converted from one format to another. This format conversion stage is an opportunity to perform transmarking that is tailored for the new format in terms of robustness and perceptibility concerns. The new format may be a broadcast format such as digital radio broadcast, or AM or FM radio broadcast. In this case, the identifier may be transmarked into a watermark or other metadata format that is robust for broadcast applications. The new format may be a compressed file format (e.g., ripping from an optical disk to an MP3 format). In this case, the identifier may be transmarked into a file header/footer or watermark format that is robust and compatible with the compressed file format.

The transmarking process may leave an existing embedded identifier in tact and layer an additional identifier into the media object. This may include encoding a new watermark that does not interfere with an existing watermark (e.g., insert the new watermark in unmarked portions of the media object or in a non-interfering transform domain). It may also include adding additional or new identifier tags to headers or footers in the file format.

Amplifying an Embedded Identifier

Rather than converting embedded data to another format, an amplifying process may be used to renew an identifier that has become weakened or separated due to processing of the media object in which it is embedded. In this case, a decoder and encoder pair may be used to determine the current identifier and re-encode it. Of course, the encoder can also choose to embed a new or additional identifiers as well.

If the previous identifier is lost, the encoder can query an identifier database established in the registration process, passing identifying information about the media object. The database uses the identifying information to find an associated identifier and returns it to the encoder for embedding in the media object.

Managing On-Line Media Library Through Links in Media Signals

The forms in which digital media content can be distributed continue to evolve rapidly. Video and audio signals can be stored in a digital content package and distributed in physical form, such as an optical or magnetic storage medium, or in an electronic form (e.g., transferred over a network in a compressed or uncompressed form). In this document, a content package refers to a format in which a title, e.g., a film, song, musical album, multimedia collection etc., is played from a complete representation of that title.

In contrast, media content may also be delivered over a wire or wireless communication link in a streaming format. Obviating the need to have a complete copy of the title, a streaming format enables the receiver to play the title as it receives portions of it in a data "stream" from an external source. The following sections describe applications for linking media signals to other content and data using metadata and/or steganography.

Linking Packaged Digital Media to On-line Library of Media Titles

In this application, a local application (e.g., a device or software process) extracts an identifier from a media signal stored in a content package, and communicates the identifier to a database application to create and manage a library of media titles. Examples of a content package include optical media such as CDs and DVDs, magnetic media such as floppy disks and tapes, flash memory, compressed media

files, etc. The user places the package into a media reader, such as a disk drive, player, etc. Operating in conjunction with the media reader, the local application extracts information (e.g., a portion of the media signal) from the package, extracts the identifier, and sends it to a database system (e.g., a server on the Internet). In response, the database system determines the corresponding title and adds the title to an on-line library (e.g., external storage accessible via the Internet). The library may be set up as a personal collection, or a collection for a group of users.

To identify the user(s)' library, the local application provides a user identifier. This user identifier may be authentication information entered by a user (such as a user name and password), or alternatively, may be an identifier (such as a device ID) sent automatically by the local application.

The title (i.e. content) is added to the on-line library, by transferring a copy of the selection (e.g., music track, video, etc.) from a master database (e.g., a library of MP3 files, or some other streaming or downloadable content format) to the user's on-line library collection. This arrangement avoids the need to upload content from the user's application. Also, it is a much more secure approach than techniques that simply read title data from a CD and relay same to the on-line library. (It is a simple task for an unscrupulous user to fake the presence of a CD by determining how the client CD software specifies the title to the on-line library, and then mimic same even without possession of a bona fide CD.) The in-band encoding presented by watermarks offers innately better security, and provides opportunities for enhanced security by encryption, etc.

In other arrangements, a copy of the selection, per se, is not transferred from the master database to the user's library, but rather a reference (e.g., a link or pointer) to the master library is added to the user's library. Efficiencies in storage can thereby be achieved (i.e., a copy of each selection is stored only once, from which an unlimited number of users' on-line libraries can link to it).

The identifier may be placed in the content package by steganographically encoding it in the media signal. For example, the identifier may be a reference number (e.g., of 24–256 bits) or the text name of the title embedded in a digital watermark. In a digital watermark implementation, a watermark embedder encodes the identifier in video, audio and/or images. The local application includes a watermark detector that reads at least a portion of the media signal from the package, detects the watermark, and reads the identifier embedded in the watermark. The detector may be implemented in a computer program (e.g., driver application, browser plug-in, etc.). A communication application, such as an Internet browser, then communicates the identifier to the database system, which may be implemented using conventional database management and Internet server software.

One advantage of this application is that it allows a user to create an on-line library of titles, and then playback those titles from the library on demand. For example, the user may organize a large collection of titles, view titles in a variety of formats, and playback individual songs or videos, in any order and at any time. The user can request playback anywhere by connecting to the on-line database and requesting a streaming delivery or file download.

For playback, a player application (e.g., device or application program on a computer) sends a request to a content delivery system via a wire or wireless connection. The content delivery system first checks to make sure that the user has the title in her on-line library. In addition, it may authenticate the user and determine usage rights before returning any content. If it determines playback to be

authorized, the content delivery system sends the titles by streaming the content to the player application, on demand, in the order requested.

Linking Streaming Media to On-line Library of Media Titles

A similar scheme to the one described in the previous section may be implemented for streaming media. In this case, the local application need not have a packaged version of the content to add a title to a user's library. Instead, the local application extracts an identifier from a portion of the streaming content. The identifier may be embedded in a watermark that is replicated throughout the media signal. In the event that the portion of the streaming media does not contain an identifier, the local application continues to execute a detection process on the media signal as it arrives until it has extracted the identifier.

In either of the above applications, the user can initiate a process of extracting the watermark by an explicit request, such as by clicking on the visual UI of the local application, entering a voice command, etc. Alternatively, the local application may initiate the detection process automatically whenever the user starts playback from packaged or streaming content.

The identifier may also include usage rights that dictate how the user (as identified by a user ID) may retrieve a copy from the library for playback. For example, the watermark may include a number that represents the number of times the user can access the content for playback.

Linking Packaged or Streaming Media to Database of Auxiliary Information Related to the Media

In addition to linking to a title database, the identifier may also link to other information or machine instructions relating to the media. For example, the database may send a set of options back to the user (e.g., in the form of a HTML page) that allow the user to select and download additional information related to the media signal in which the identifier is embedded.

Operating Environment for Computer Implementations

FIG. 3 illustrates an example of a computer system that serves as an operating environment for software implementations of the systems described above. The software applications may be implemented in C/C++ and are portable to many different computer systems. FIG. 3 generally depicts one such system.

The computer system shown in FIG. 3 includes a computer 1220, including a processing unit 1221, a system memory 1222, and a system bus 1223 that interconnects various system components including the system memory to the processing unit 1221.

The system bus may comprise any of several types of bus structures including a memory bus or memory controller, a peripheral bus, and a local bus using a bus architecture such as PCI, ISA and EISA, to name a few.

The system memory includes read only memory (ROM) 1224 and random access memory (RAM) 1225. A basic input/output system 1226 (BIOS), containing the basic routines that help to transfer information between elements within the computer 1220, such as during start-up, is stored in ROM 1224.

The computer 1220 further includes a hard disk drive 1227, a magnetic disk drive 1228, e.g., to read from or write to a removable disk 1229, and an optical disk drive 1230, e.g., for reading a CD-ROM or DVD disk 1231 or to read from or write to other optical media. The hard disk drive 1227, magnetic disk drive 1228, and optical disk drive 1230 are connected to the system bus 1223 by a hard disk drive interface 1232, a magnetic disk drive interface 1233, and an optical drive interface 1234, respectively. The drives and

their associated computer-readable media provide nonvolatile storage of data, data structures, computer-executable instructions (program code such as dynamic link libraries, and executable files), etc. for the computer 1220.

Although the description of computer-readable media above refers to a hard disk, a removable magnetic disk and an optical disk, it can also include other types of media that are readable by a computer, such as magnetic cassettes, flash memory cards, digital video disks, and the like.

A number of program modules may be stored in the drives and RAM 1225, including an operating system 1235, one or more application programs 1236, other program modules 1237, and program data 1238.

A user may enter commands and information into the personal computer 1220 through a keyboard 1240 and pointing device, such as a mouse 1242. Other input devices may include a microphone, joystick, game pad, satellite dish, digital camera, scanner, or the like. The microphone may be used to capture audio signals. Similarly, a digital camera or scanner 43 may be used to capture video and images. The camera and scanner are each connected to the computer via a standard interface 44. Currently, there are digital cameras designed to interface with a Universal Serial Bus (USB), Peripheral Component Interconnect (PCI), and parallel port interface. Two emerging standard peripheral interfaces for cameras include USB2 and 1394 (also known as firewire and iLink).

These and other input devices are often connected to the processing unit 1221 through a serial port interface 1246 that is coupled to the system bus, but may be connected by other interfaces, such as a parallel port, game port or a universal serial bus (USB).

A monitor 1247 or other type of display device is also connected to the system bus 1223 via an interface, such as a video adapter 1248. In addition to the monitor, personal computers typically include other peripheral output devices (not shown), such as speakers and printers.

The computer 1220 operates in a networked environment using logical connections to one or more remote computers, such as a remote computer 1249. The remote computer 1249 may be a server, a router, a peer device or other common network node, and typically includes many or all of the elements described relative to the computer 1220, although only a memory storage device 1250 has been illustrated in FIG. 3. The logical connections depicted in FIG. 3 include a local area network (LAN) 1251 and a wide area network (WAN) 1252. Such networking environments are commonplace in offices, enterprise-wide computer networks, intranets and the Internet.

When used in a LAN networking environment, the computer 1220 is connected to the local network 1251 through a network interface or adapter 1253. When used in a WAN networking environment, the personal computer 1220 typically includes a modem 1254 or other means for establishing communications over the wide area network 1252, such as the Internet. The modem 1254, which may be internal or external, is connected to the system bus 1223 via the serial port interface 1246.

In a networked environment, program modules depicted relative to the personal computer 1220, or portions of them, may be stored in the remote memory storage device. The processes detailed above can be implemented in a distributed fashion, and as parallel processes. It will be appreciated that the network connections shown are exemplary and that other means of establishing a communications link between the computers may be used.

The computer may establish a wireless connection with external devices through a variety of peripherals such as a cellular modem, radio transceiver, infrared transceiver, etc.

While a computer system is offered as example operating environment, the applications may be implemented in a variety of devices and systems, including servers, workstations, hand-held devices (e.g., hand held audio or video players, Personal Digital Assistants such as Palm Pilot, etc.), network appliances, distributed network systems, etc.

Concluding Remarks

Having described and illustrated the principles of the technology with reference to specific implementations, it will be recognized that the technology can be implemented in many other, different, forms. To provide a comprehensive disclosure without unduly lengthening the specification, applicants incorporate by reference the patents and patent applications referenced above. These patents and patent applications provide additional implementation details. They describe ways to implement processes and components of the systems described above. Processes and components described in these applications may be used in various combinations, and in some cases, interchangeably with processes and components described above.

The particular combinations of elements and features in the above-detailed embodiments are exemplary only; the interchanging and substitution of these teachings with other teachings in this and the incorporated-by-reference patents/applications are also contemplated.

We claim:

1. A method of establishing an on-line collection of media titles, comprising the steps of:

30 extracting an identifier steganographically encoded in a media signal;
sending the identifier to a database; and
requesting the database to add a title associated with the media signal to an on-line collection.

2. The method according to claim 1, further comprising the step of:

requesting playback of a title from the on-line collection.

3. A system to interact with an on-line collection of media titles stored in a database, said system comprising:

40 a detector to extract an embedded watermark from a media signal, the watermark including identifying information for the media signal;
a module to send the identifying information to the database; and
a module to receive data from the database in response to the identifying information.

45 4. The system according to claim 3, wherein the identifying information includes a title of the media signal to be added to the on-line collection.

50 5. The system according to claim 3, wherein the identifying information includes usage rights.

6. The system according to claim 3, wherein the media signal is stored in a content package.

55 7. An apparatus to interact with a media signal, said apparatus comprising:

60 a storage device;
a processing unit; and
a local application stored in said storage device and processed by said processing unit, said local application operating to: i) extract an identifier from the media signal, the identifier being steganographically encoded into the media signal, and to ii) send the identifier to a database to associate the media signal with an on-line collection maintained by the database.

65 8. The apparatus according to claim 7, wherein the identifier comprises a watermark.

21

- 9. The apparatus according to claim 7, wherein the identifier comprises a reference number.
- 10. The apparatus according to claim 7, wherein the identifier comprises data embedded within a watermark.
- 11. The apparatus according to claim 10, wherein the data comprises a title name.
- 12. The apparatus according to claim 7, wherein said local application provides a user identifier to the database.
- 13. The apparatus according to claim 12, wherein said user identifier comprises a device identifier.
- 14. The apparatus according to claim 7, wherein the media signal is stored in a content package.
- 15. A method of operating a system to create and manage a library of on-line media titles, said method comprising the steps of:
 - receiving a media identifier from a user that has been extracted from steganographically encoded information in a media signal;
 - determining a title that corresponds to the media identifier;
 - identifying the user; and
 - adding the media title to an on-line library of media titles associated with the user.
- 16. The method according to claim 15, further comprising the step of:
 - in response to a user's request for a media title, verifying that the media title is in the user's on-line library.
- 17. The method according to claim 16, further comprising the steps of:
 - authenticating the user in response to the user's request for a media title; and
 - allowing the user to access the on-line library when the user is authenticated.
- 18. The method according to claim 17, further comprising the steps of:
 - determining usage rights associated with a requested media title, and controlling access to the media title based on the usage rights.
- 19. The method according to claim 15, wherein said adding step comprises the step of transferring a copy of the media title to the on-line library.

22

- 20. The method according to claim 15, wherein said adding step comprises the step of adding a pointer to the on-line library to point to a copy of the media title.
- 21. The method according to claim 15, wherein the user is identified by identifying a user device in communication with the system.
- 22. An apparatus to interact with a streaming media signal, said apparatus comprising:
 - a storage device;
 - a processing unit; and
 - a local application stored in said storage device and processed by said processing unit, said local application operating to:
 - i) extract an identifier from the streaming media signal, and to
 - ii) send the identifier to a database to associate the streaming media signal with an on-line library collection maintained by the database.
- 23. The apparatus according to claim 22, wherein the identifier is embedded in a watermark in the streaming media signal.
- 24. The apparatus according to claim 23, wherein the watermark is replicated throughout the streaming media signal.
- 25. The apparatus according to claim 23, wherein the local application extracts the identifier upon a request by a user.
- 26. The apparatus according to claim 25, wherein the request comprises a voice command.
- 27. The apparatus according to claim 23, wherein the identifier comprises usage rights which dictate how the user may retrieve a copy of the streaming media signal from the library for playback.
- 28. The apparatus according to claim 23, wherein the on-line library is a personal library associated with a pre-determined user.
- 29. The apparatus according to claim 23, wherein the on-line library is associated with a predetermined group of users.

* * * * *



US005774670A

United States Patent [19]

Montulli

[11] Patent Number: 5,774,670
[45] Date of Patent: Jun. 30, 1998

[54] PERSISTENT CLIENT STATE IN A HYPERTEXT TRANSFER PROTOCOL BASED CLIENT-SERVER SYSTEM

[75] Inventor: Lou Montulli, Palo Alto, Calif.

[73] Assignee: Netscape Communications Corporation, Mountain View, Calif.

[21] Appl. No.: 540,342

[22] Filed: Oct. 6, 1995

[51] Int. Cl.⁶ G06F 13/38; G06F 17/30

[52] U.S. Cl. 395/200.57; 395/200.33; 395/200.47; 345/335

[58] Field of Search 395/200.32, 200.33, 395/300.48, 200.49, 200.57, 200.58, 200.6; 345/335; 707/501, 10

[56] References Cited

PUBLICATIONS

Van Name et al., "Putting the lid on Pandora's Cookie Jar"; PC Week Aug. 19, 1996 V13 N33 PN8t1).

Foster, "Can mixing 'cookies' with online marketing be a recipe for heartburn?"; InfoWorld Jul. 22, 1996 v8 p54(1).

McCarthy, "The Netscape Biscuit Company", Government Computer News Sep. 23, 1996 v15 no 24 p55(2).

Raynovich, "Microsoft readies browser update"; LAN Times; Aug. 5, 1996 v13 no 17 p7(1).

Netscape; "Persistent Client State HTTP Cookies"; 1997; http://home.netscape.com./newsref/std/cookie_spec.htm.

Montulli et al.; "Proposed HTTP State Management Mechanism"; Aug. 16, 1996 HTTP Working Group.

"Cookie I-D Drafts"; Nov. 21, 1997; http://portal.research.bell-labs.com/dmk/cookie-ver.html.

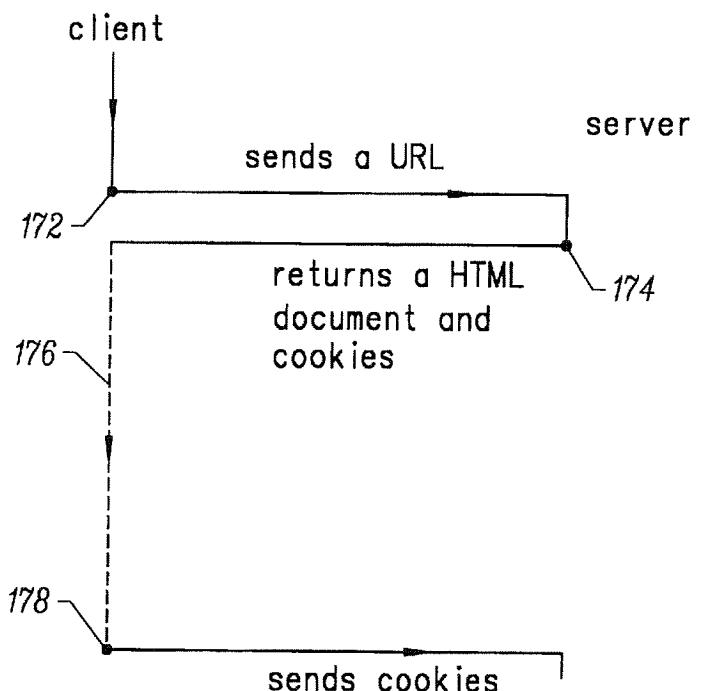
CMP NetGuide Magazine; "The Trouble With Cookies", May 1, 1996; Issue 305.

Primary Examiner—Mark H. Rinehart
Attorney, Agent, or Firm—Michael A. Glenn; Harvey J. Anderson

[57] ABSTRACT

A method and apparatus for transferring state information between a server computer system and a client computer system. In one embodiment of the method, an http client requests a file, such as an HTML document, on an http server, and the http server transmits the file to the http client. In addition, the http server transmits a state object, which describes certain state information, to the http client. The http client stores the state object, and will typically send the state object back to the http server when making later requests for files on the http server. In a typical embodiment, the state object includes a domain attribute which specifies a domain or network address, and the state object is transmitted from the http client to a server only when the http client makes an http request to the server and the server is within the domain. In one embodiment, the apparatus includes a processor and memory and a computer readable medium which stores program instructions. In the case of the client system, the instructions specify operations such as receiving and storing the state information; in the case of the server system, the instructions specify operations such as sending the state information to a client system.

26 Claims, 8 Drawing Sheets



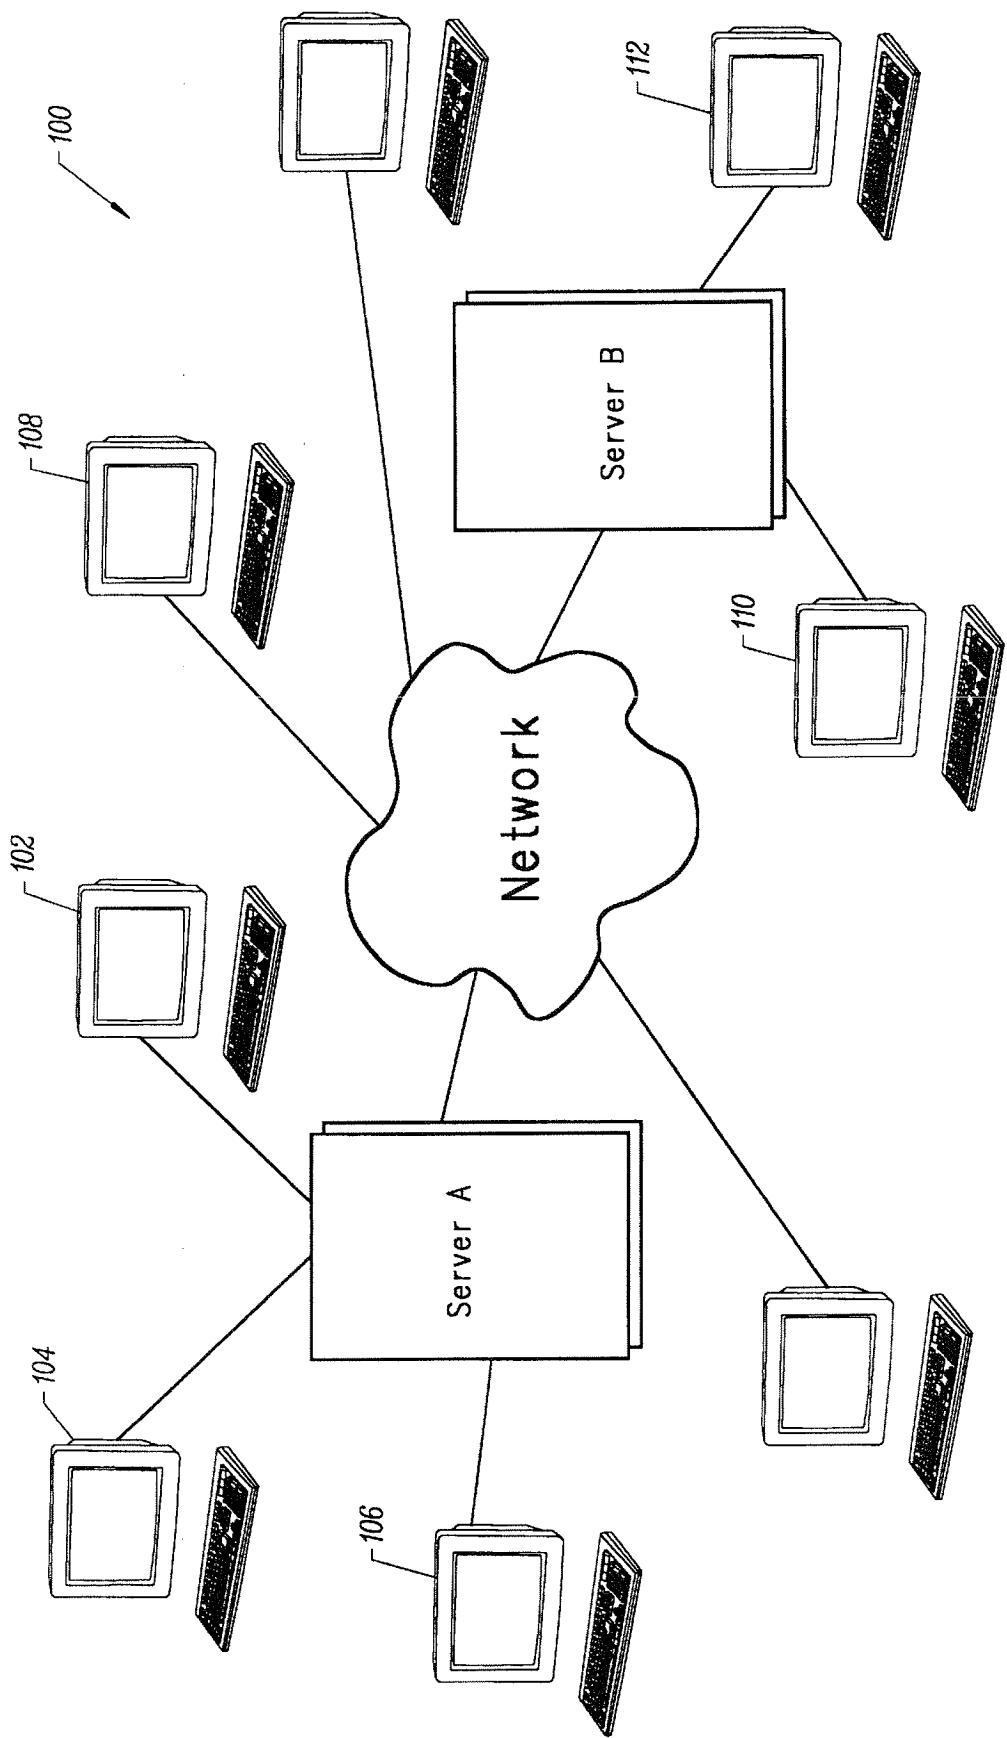


FIG. 1A

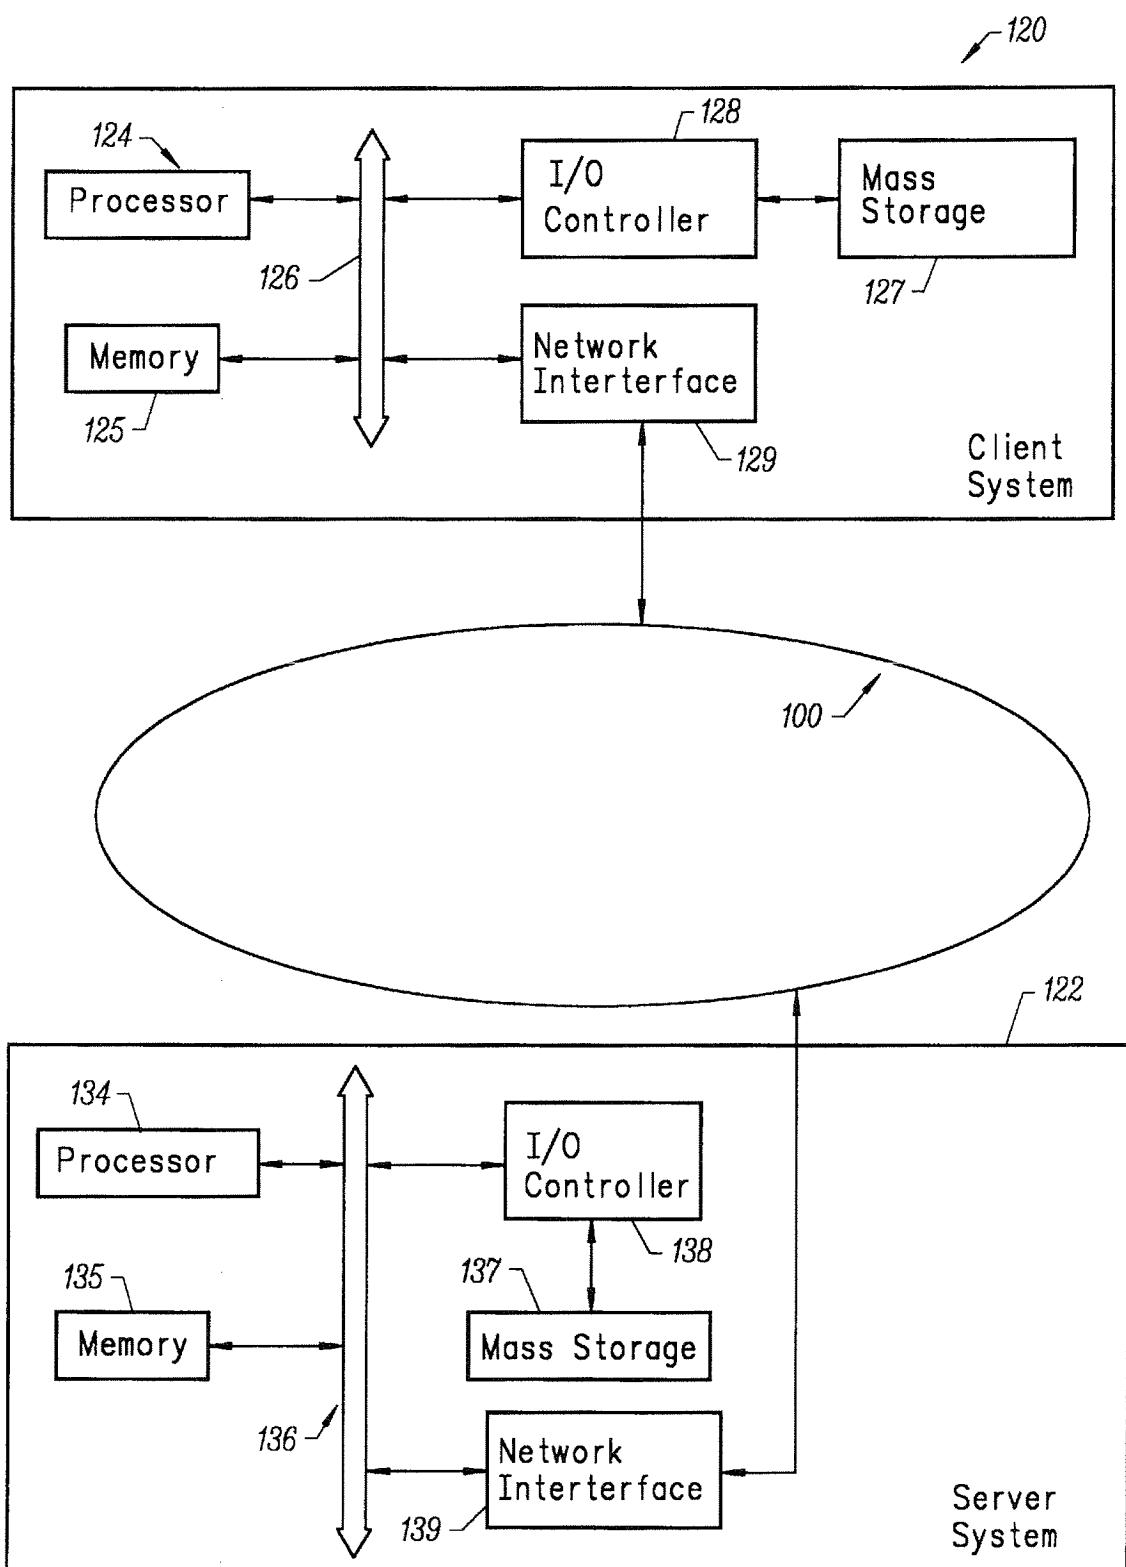


FIG. 1B

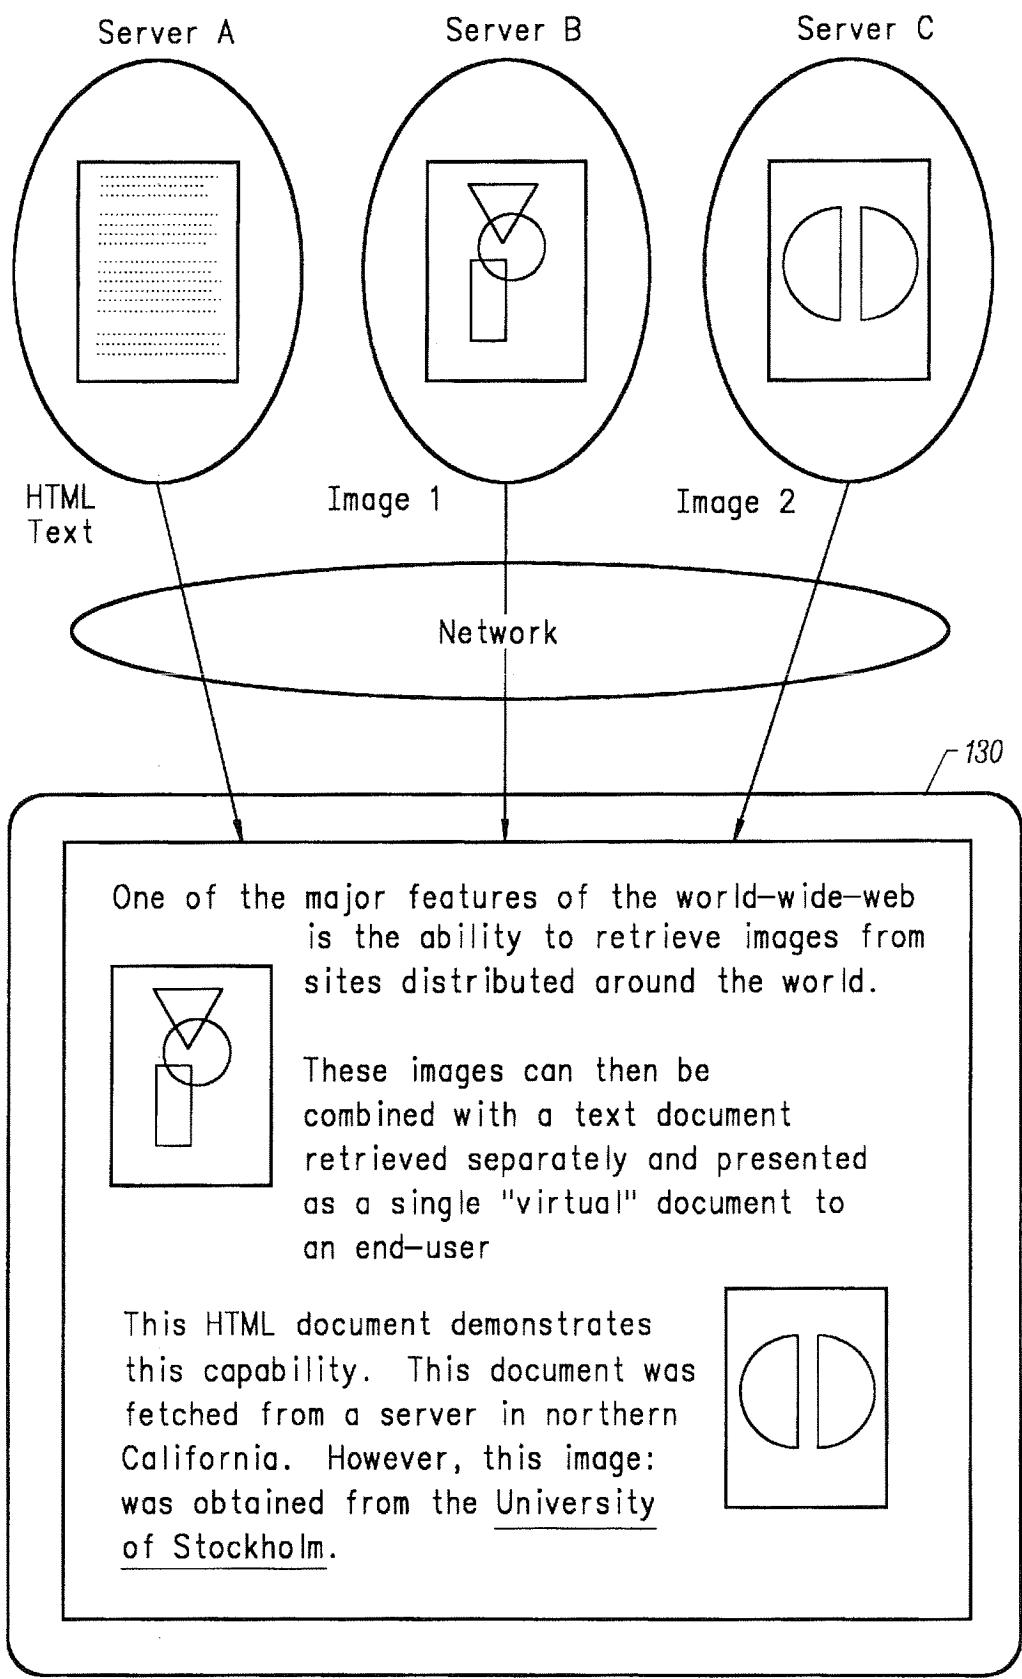


FIG. 2

```
<Title>Distributed Image Loading Example </Title>
<h1>Distributed Image Loading Example</h1>
```

One of the major features of the world-wide-web is the ability to retrieve images from sites distributed around the world. These images can then be combined with a text document retrieved separately and presented as a single "virtual" document to an end-user. This HTML document demonstrates this capability.<p>

This document was fetched from a server in northern California.
However, this image:

was obtained from Illinois.<P>

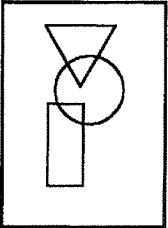
And this one: came
from the University of
Stockholm.

FIG. 3A

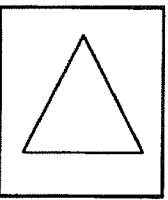
Distributed Image Loading Example

One of the major features of the world-wide web is the ability to retrieve images from sites distributed around the world. These images can then be combined with a text document retrieved separately and presented as a single "virtual" document to an end-user. This HTML document demonstrates this capability.

This document was fetched from a server in Northern California. However, this image:



was obtained from Illinois. And this one:



was obtained from the University of Stockholm.

FIG. 3B

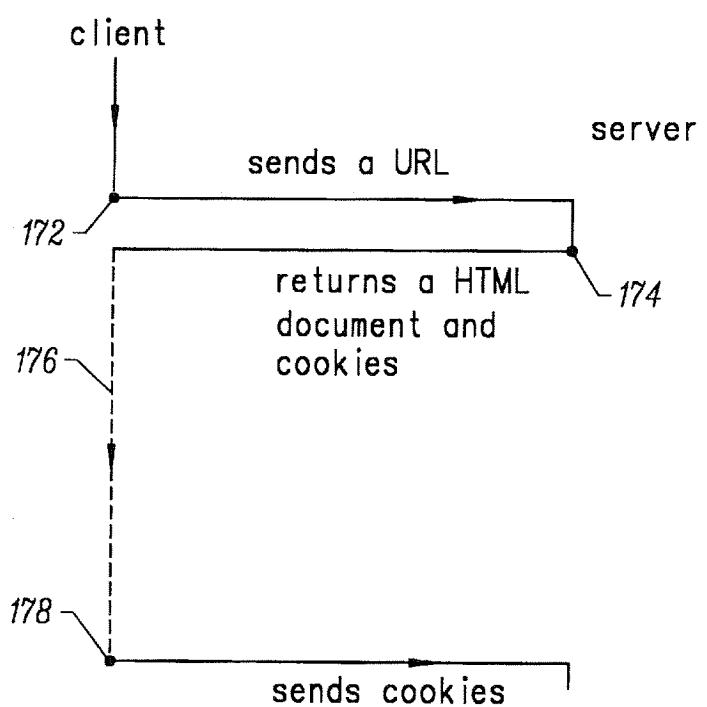


FIG. 4

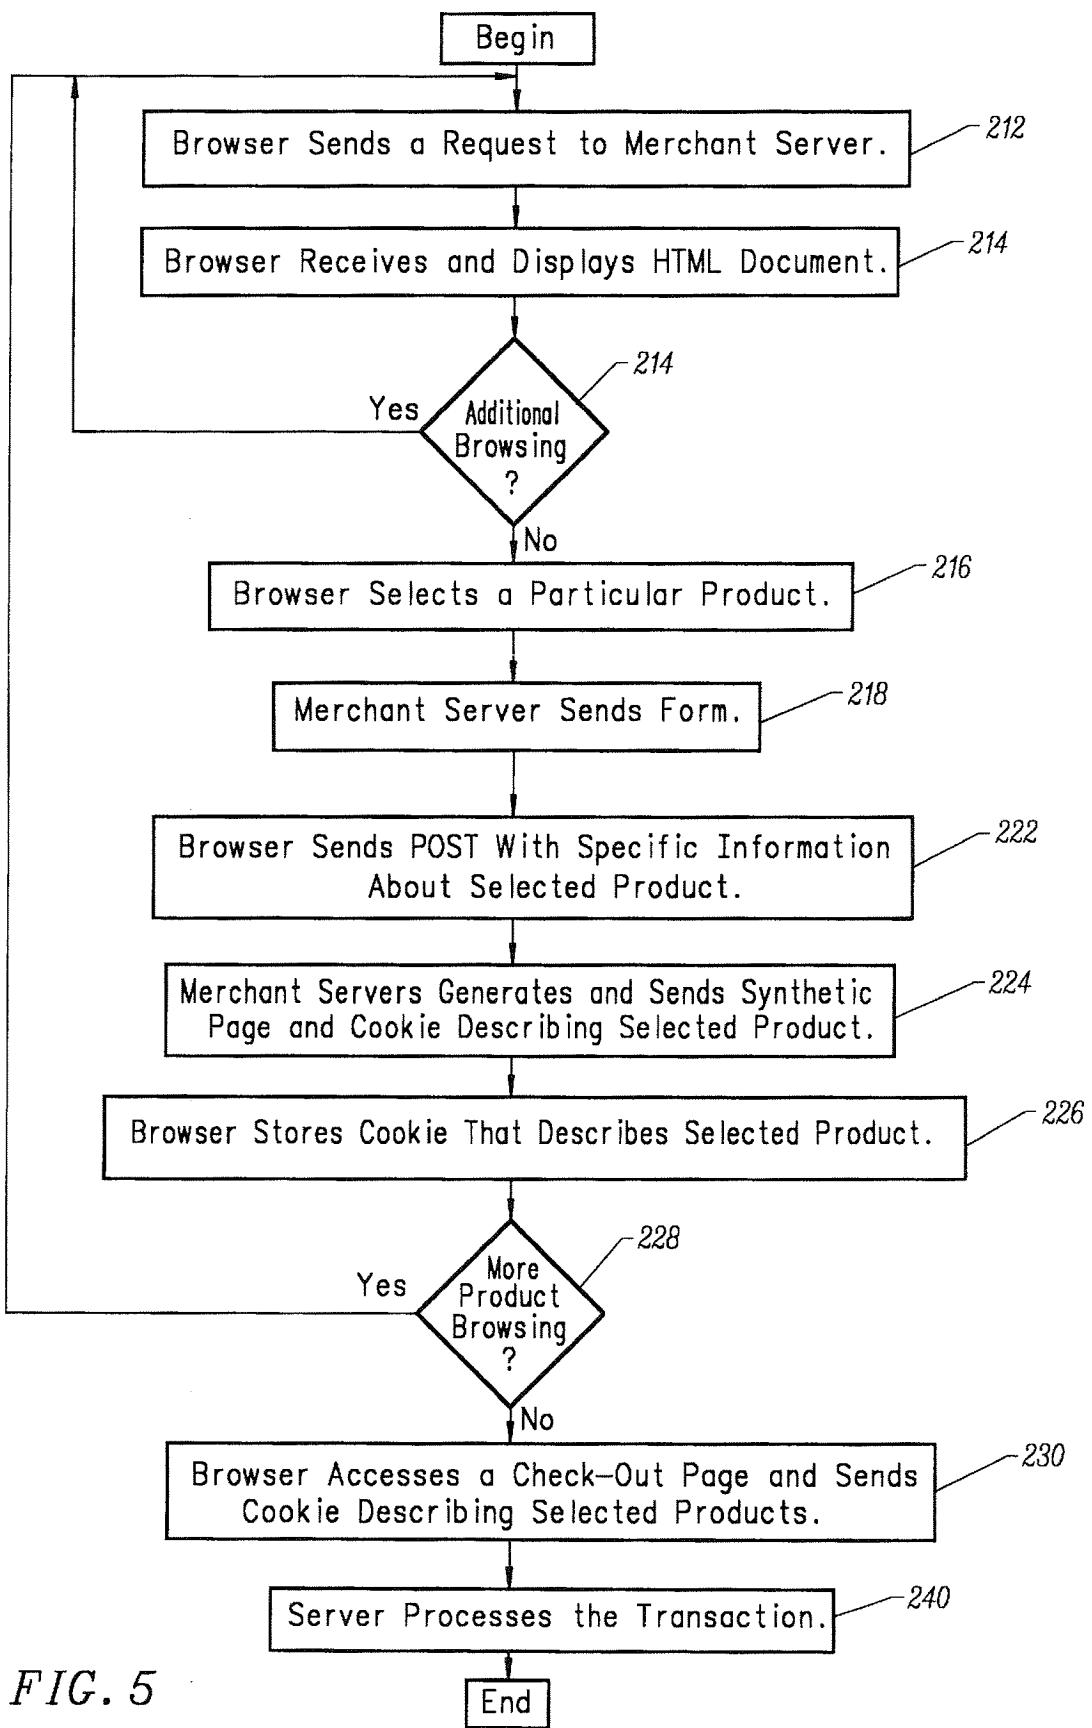


FIG. 5

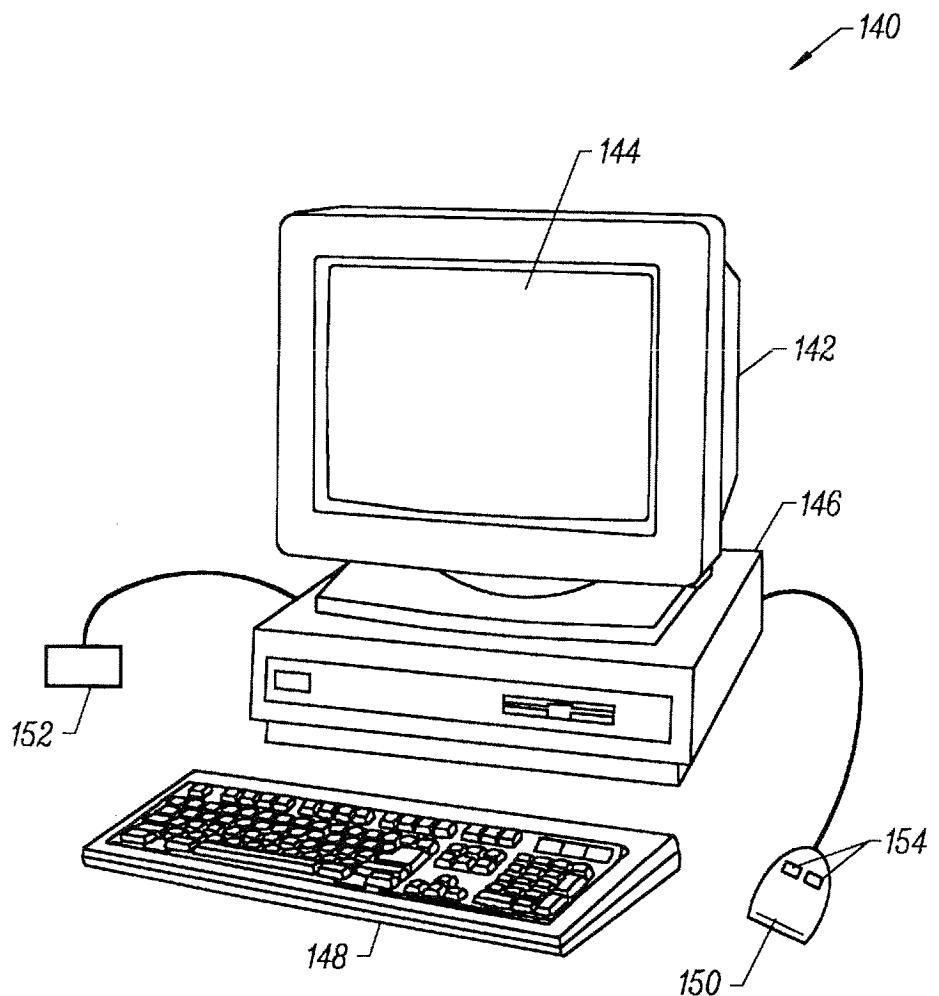


FIG. 6

**PERSISTENT CLIENT STATE IN A
HYPertext TRANSFER PROTOCOL
BASED CLIENT-SERVER SYSTEM**

FIELD OF THE INVENTION

This invention relates to communication in a client-server computer systems. Specifically, the invention relates to client-server computer systems in which a server can send state information to a client and the client stores the state information for later retransmissions back to the server.

BACKGROUND OF THE INVENTION

An important use of computers is the transfer of information over a network. Currently, the largest computer network in existence is the InterNet. The InterNet is a worldwide interconnection of computer networks that communicate using a common protocol. Millions of computers, from low end personal computers to high-end super computers are coupled to the InterNet.

The InterNet grew out of work funded in the 1960s by the U.S. Defense Department's Advanced Research Projects Agency. For a long time, InterNet was used by researchers in universities and national laboratories to share information. As the existence of the InterNet became more widely known, many users outside of the academic/research community (e.g., employees of large corporations) started to use InterNet to carry electronic mail.

In 1989, a new type of information system known as the World-Wide-Web ("the Web") was introduced to the InterNet. Early development of the Web took place at CERN, the European Particle Physics Laboratory. The Web is a wide-area hypermedia information retrieval system aimed to give wide access to a large universe of documents. At that time, the Web was known to and used by the academic/research community only. There was no easily available tool which allows a technically untrained person to access the Web.

In 1993, researchers at the National Center for Supercomputing Applications (NCSA) released a Web browser called "Mosaic" that implemented a graphical user interface (GUI). Mosaic's graphical user interface was simple to learn yet powerful. The Mosaic browser allows a user to retrieve documents from the World-Wide-Web using simple point-and-click commands. Because the user does not have to be technically trained and the browser is pleasant to use, it has the potential of opening up the InterNet to the masses.

The architecture of the Web follows a conventional client-server model. The terms "client" and "server" are used to refer to a computer's general role as a requester of data (the client) or provider of data (the server). Under the Web environment, Web browsers reside in clients and Web documents reside in servers. Web clients and Web servers communicate using a protocol called "HyperText Transfer Protocol" (HTTP). A browser opens a connection to a server and initiates a request for a document. The server delivers the requested document, typically in the form of a text document coded in a standard Hypertext Markup Language (HTML) format, and when the connection is closed in the above interaction, the server serves a passive role, i.e., it accepts commands from the client and cannot request the client to perform any action.

The communication model under the conventional Web environment provides a very limited level of interaction between clients and servers. In many systems, increasing the level of interaction between components in the systems often makes the systems more robust, but increasing the

interaction increases the complexity of the interaction and typically slows the rate of the interaction. Thus, the conventional Web environment provides less complex, faster interactions because of the Web's level of interaction between clients and servers.

In the conventional Web environment, clients do not retain information of a session after the session is closed. In many systems, the ability to retain information after the systems become inactive is crucial to the functioning of the systems. Thus, it is desirable to allow clients to have this ability.

SUMMARY OF THE INVENTION

The present invention involves a client-server system on a network in which a server can send state information to a client and the client stores the state information. The stored state information can later be sent back to the server at appropriate times. In this manner, the state of a client can be maintained in the client-server system where no state inherently exists in such a system.

One embodiment of the present invention is a network system for communicating documents containing information such as text and one or more images. The system comprises a first computer (i.e., a server) capable of sending such documents over a network such as the InterNet. The system also has a second computer (i.e., a client) which can request these documents or files from the server. After the requested documents are received, the client can display the documents. In accordance with the present invention, the server can send state information to the client when a document is sent. The client then stores the state information, which is typically in the form of a state object. In a subsequent request for documents to the server, the client can send the stored state information to the server.

In an embodiment of the invention, the server uses a hypertext transfer protocol ("http") to communicate over the network with clients; such clients also communicate with the server using the hypertext transfer protocol. This server and these clients are referred to as an http server and http clients respectively. The server typically will include a server processor and a memory and a computer readable medium, such as a magnetic ("hard disk") or optical mass storage device, and the computer readable medium of the server contains computer program instructions for transmitting the file from the server system to the client system and for transmitting the state object to the client system. The client typically will include a client processor and a memory and a computer readable medium, such as a magnetic or optical mass storage device, and the computer readable medium of the client contains computer program instructions for receiving the state object, which specifies the state information, from the server and for storing the state object at the client. The state object, in a typical embodiment, will include a name attribute, such as a domain attribute.

One of the applications of the present invention is an on-line shopping system. A customer can browse information delivered by a merchant server using a browser running on a client. The customer can also select products to be placed in a virtual shopping basket. The server then sends state information related to the selected products to the browser on the client for storage. When the customer wants to purchase the products in the virtual shopping basket, the browser sends the corresponding state information to a specified check-out Web page for processing.

Another application of the present invention is an "on-line" information service, such as a newspaper's Web server

which includes articles or other information from the newspaper's subscription services. In one example, a newspaper or publishing company may have several different publications, each requiring a separate subscription fee which may differ among the different publications. A user of the information service may browse the different publications by making http requests, from the client's/user's computer system, to the publisher's Web server which responds with the requested publication and state information specifying the user's identification, and other subscription information (e.g., user registration and billing information) which allows the user to view the contents of the publication; this information is typically provided by the user at least once in a conventional log-on process. Thereafter, this information is included in the state information which is exchanged between the client and the server in the process of the invention. Accordingly, when the user, during the browsing process, desires to view another publication (e.g., from the same or different publisher) this state information will be transmitted back to the Web server to provide the necessary subscription information (thereby entitling the user to view the publication) without requiring the user to re-enter the necessary subscription information. In this manner, a user may browse from publication to publication on the Web server or a different Web server in the domain without having to re-enter, when seeking a new publication, the necessary subscription information.

These and other features of the present invention will be disclosed in the following description of the invention together with the accompanying drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

The objects, features, and advantages of the present invention will be apparent from the following detailed description of the preferred embodiment of the invention with references to the following drawings.

FIG. 1A is a pictorial diagram of a computer network used in the present invention.

FIG. 1B shows a computer network containing a client system and a server system.

FIG. 2 illustrates the retrieval of remote text and images and their integration in a document.

FIG. 3A shows an example of an HTML document which can be processed by the browser of the present invention.

FIG. 3B shows the integrated document corresponding to the HTML document of FIG. 3A as it would appear on a display screen of a client computer.

FIG. 4 shows schematically the flow of information between a client and a server in accordance with the present invention.

FIG. 5 is a flow chart showing the operation of a merchant system of the present invention.

FIG. 6 shows a computer system that could be used to run the browser of the present invention.

DETAILED DESCRIPTION

Methods and apparatuses for maintaining state information in a client-server based computer network system are disclosed. The following description is presented to enable any person skilled in the art to make and use the invention. For purposes of explanation, specific nomenclature is set forth to provide a thorough understanding of the present invention. However, it will be apparent to one skilled in the art that these specific details are not required to practice the present invention. Descriptions of specific applications are

provided only as examples. Various modifications to the preferred embodiments will be readily apparent to those skilled in the art, and the general principles defined herein may be applied to other embodiments and applications without departing from the spirit and scope of the invention. Thus, the present invention is not intended to be limited to the embodiments shown, but is to be accorded the widest scope consistent with the principles and features disclosed herein.

Prior to describing the present invention, some introductory material is explained, including explanations concerning client-server computing, InterNet addresses, URL's and browsing of the Web.

CLIENT-SERVER COMPUTING

FIG. 1A illustrates a conceptual diagram of a computer network 100, such as the InterNet. Computer network 100 comprises small computers (such as computers 102, 104, 106, 108, 110 and 112) and large computers, such as computers A and B, commonly used as servers. In general, small computers are "personal computers" or workstations and are the sites at which a human user operates the computer to make requests for data from other computers or servers on the network. Usually, the requested data resides in large computers. In this scenario, small computers are clients and the large computers are servers. In this specification, the terms "client" and "server" are used to refer to a computer's general role as a requester of data (client) or provider of data (server). In general, the size of a computer or the resources associated with it do not preclude the computer's ability to act as a client or a server. Further, each computer may request data in one transaction and provide data in another transaction, thus changing the computer's role from client to server, or vice versa.

A client, such as computer 102, may request a file from server A. Since computer 102 is directly connected to server A through a local area network, this request would not normally result in a transfer of data over what is shown as the "network" of FIG. 1. The "network" of FIG. 1 represents the InterNet which is an interconnection of networks. A different request from computer 102 may be for a file that resides in server B. In this case, the data is transferred from server B through the network to server A and, finally, to computer 102. The distance between servers A and B may be very long, e.g. across continents, or very short, e.g., within the same city. Further, in traversing the network the data may be transferred through several intermediate servers and many routing devices, such as bridges and routers.

The World-Wide-Web ("The Web") uses the client-server model to communicate information between clients and servers. Web Servers are coupled to the InterNet and respond to document requests from Web clients. Web clients (also known as Web "browsers") are programs that allow a user to simply access Web documents located on Web Servers.

FIG. 1B shows, in more detail, an example of a client-server system interconnected through the InterNet 100. In this example, a remote server system 122 is interconnected through the InterNet to client system 120. The client system 120 includes conventional components such as a processor 124, memory 125 (e.g. RAM), a bus 126 which couples the processor 124 and memory 125, a mass storage device 127 (e.g. a magnetic hard disk or an optical storage disk) coupled to the processor and memory through an I/O controller 128 and a network interface 129, such as a conventional modem. The server system 122 also includes conventional components such as a processor 134, memory 135 (e.g. RAM), a

bus 136 which couples the processor 134 and memory 135, a mass storage device 137 (e.g. a magnetic or optical disk) coupled to the processor 134 and memory 135 through an I/O controller 138 and a network interface 139, such as a conventional modem. It will be appreciated from the description below that the present invention may be implemented in software which is stored as executable instructions on a computer readable medium on the client and server systems, such as mass storage devices 127 and 137 respectively, or in memories 125 and 135 respectively.

InterNet Addresses

As discussed in the background, the InterNet consists of a worldwide computer network that communicates using well defined protocol known as the InterNet Protocol (IP). Computer systems that are directly connected to the InterNet each have an unique InterNet address. An InterNet address consists of four numbers where each number is less than 256. The four numbers of an InterNet address are commonly written out separated by periods such as 192.101.0.3

To simplify InterNet addressing, the "Domain Name System" was created. The domain name system allows users to access InterNet resources with a simpler alphanumeric naming system. An InterNet Domain name consists of a series of alphanumeric names separated by periods. For example, the name "drizzle.stanford.edu" is the name for a computer in the physics department at Stanford University. Read from left to right, each name defines a subset of the name immediately to the right. In this example, "drizzle" is the name of a workstation in the "stanford" domain. Furthermore, "stanford" is a subset of the "edu" domain. When a domain name is used, the computer accesses a "Domain Name Server" to obtain the explicit four number InterNet address.

Uniform Resource Locators

To further define the addresses of resources on the InterNet, the Uniform Resource Locator system was created. A Uniform Resource Locator (URL) is a descriptor that specifically defines a type of InterNet resource and its location. URLs have the following format:

resource_type: / / domain.address/ path_name Where
 "resource_type" defines the type of internet resource.
 Web documents are identified by the resource type "http" which indicates that the hypertext transfer protocol should be used to access the document. Other resource types include "ftp" (file transmission protocol) and "telnet". The "domain.address" defines the domain name address of the computer that the resource is located on. Finally, the "path_name" defines a directory path within the file system of the server that identifies the resource. The right most name on the path name portion is usually the name of an actual file. Web pages are identified by the resource type "http". By convention, most Web pages end with the suffix ".html" that suggests the file is a HyperText Markup Language document. An example of a URL for a Web document is:

```
http: // info.cern.ch/hypertext/Datasources /WWW/  

Geographical.html
```

This URL indicates that by using the HTTP (Web) protocol to reach a server called "info.cern.ch", there is a directory "hypertext/Datasources/WWW" that contains a hypertext document named "Geographical.html." Resources on the Internet are uniquely addressable by their URL.

Browsing the World-Wide-Web

To access an initial Web document, the user enters the URL for a Web document into a Web browser program. The

Web browser then sends an http request to the server that has the Web document using the URL. The Web server responds to the http request by sending the requested HTTP object to the client. In most cases, the HTTP object is an plain text (ASCII) document containing text (in ASCII) that is written in HyperText Markup Language (HTML). The HTML document usually contains hyperlinks to other Web documents. The Web browser displays the HTML document on the screen for the user and the hyperlinks to other Web documents are emphasized in some fashion such that the user can selected the hyperlink.

FIG. 2 illustrates the retrieval of remote text and images and their integration in a Web page by a client computer 130. In FIG. 2, server A contains a text document coded in a standard HTML format. Server B contains an image file called Image 1 and server C contains another image file called Image 2. Each of these servers is remotely located from the other servers and the client 130. The transfer of data is via the Internet. It should be appreciated that the text and image files could be located in the same server which is remote from client 130.

FIG. 3A shows an example of an HTML document. FIG. 3B shows the corresponding integrated document (Web page) as it would appear on a display screen of a client computer. The first line of the document in FIG. 3A reads "<title>Distributed Image Loading Example</title>." In this case, the tags <title> and </title> are HTML delimiters corresponding to the beginning and ending, respectively, of text that is designated as the title of the HTML document. The title could be used for various purposes, such as listing of the document in an automatically generated index.

The second line of the HTML document of FIG. 3A reads "<h1> Distributed Image Loading Example </h1>" The <h1> and </h1> are HTML delimiters for a header that is to be displayed in a largest font. The browser software running on the client computer 130 interprets the header tags and thus displays the text between the header tags in a largest font size on the client's display screen.

After the title and header, the HTML document of FIG. 3A contains the text "One of the major features . . . capability". At the end of the text paragraph is another HTML tag shown as <p>. This is a tag indicating the end of a paragraph.

To continue with the second paragraph of the HTML document, the text reads "This document . . . this image: was obtained . . . ". The text in angle brackets defines an image to be placed in the text. Specifically, the "IMG" tag indicates that an image is being defined. The "align=middle" tag indicates that the image should be aligned in the middle of the current line of text. Finally, "src=" tag indicates that the source image file can be located using the URL "http: / /www.su.se/ SUlogo.gif".

The line continues with the phrase "from the University of Stockholm" This phrase defines "University of Stockholm" as a link to another Web document. Specifically, the "A" tag defines the beginning of a link. The "href=" tag defines that the link is to a Web page that can be located using the URL "http://www.su.se/index.html" Next, the text "University of Stockholm" is the text that will be the link. Finally, the "/A" tag defines the end of the link definition. As illustrated in FIG. 3B, the text "University of Stockholm" is displayed with underlining that indicates it is a link to another document. If the user selects the underlined text "University of Stockholm", then the browser will send out a http request for the Web page at the URL address "http: / /www.su.se/index.html".

It can be seen from the above example that the HTML document contains all information a browser needs for displaying a Web page. Thus, the only responsibility of a Web server is to provide the requested document, and there is no need for the server to request a client to do anything else. However, this role of a server also limits the utility of the Web environment.

ADDING STATE INFORMATION TO THE HYPertext TRANSFER PROTOCOL

The present invention provides an extension to the prior art HTTP protocol. Using the teachings of the present invention, when a server responds to an http request by returning an HTTP object to a client, the server may also send a piece of state information that the client system will store. In an embodiment of the present invention, the state information is referred to as a "cookie". Included in the state information (the cookie) is a description of a range of URLs for which that state information should be repeated back to. Thus, when the client system sends future HTTP requests to servers that fall within the range of defined URLs, the requests will include a transmittal of the current value of the state object. By adding the ability to transfer state information back and forth, Web servers can then play an active role in transactions between clients and servers. The term state object is also used herein to refer to the state information.

FIG. 4 is a drawing showing schematically the flow of information between a client system and a server system. At a time indicated by numeral 172, the client system sends an http request to the Web server. In response to the http request, the server returns an HTML document together with a header, which is typically separate from the HTML documents, at a time indicated by numeral 174. The header may contain one or more cookies. Upon receiving the HTML document and the header, the client system displays an HTML document on the screen and stores the cookies in a memory such as a storage medium. The client may switch and work on other tasks, or may be shut down completely. This is indicated by a dash line 176. At a time indicated by numeral 178, the client system may access a Web server that is specified in the received cookie such that the client system transmits the cookies to the server, thus providing state information about the client system to the server system.

This extension to the http protocol provides a powerful new tool which enables a large number of new types of applications to be written for a Web-based environment. Examples of new applications include on-line shopping that stores information about items currently selected by consumers, for-fee on-line services that can send back registration information and thus free users from retyping a user-id on next connection, and Web sites that can store per-user preferences on the client system and have the client supply those preferences every time the site is later accessed.

Server Behavior

A particular embodiment of the state information is described below in order to provide an example according to the present invention. It will be appreciated that alternative formats may be used in accordance with the principles of the present invention. As stated above, the extension to the HTTP protocol adds a new piece of state information to the HTTP header as part of an HTTP response from a Web server. Typically, the state information is generated by a common gateway interface ("CGI") script. The state information is stored by the receiving client system in the form of a "cookie list" for later use. The syntax of the new data, in one embodiment, is:

Set-Cookie: NAME=VALUE; expires=DATE; path=
PATH;

domain=DOMAIN_NAME; secure

The capitalized terms can be set by the server system. The first attribute is "NAME=VALUE". This attribute serves to identify a cookie. The "NAME" attribute is a name for the cookie. The "NAME" attribute is the only required attribute on the "Set-Cookie" header in one embodiment. The "VALUE" is a value assigned to the previously defined name. The "VALUE" is a string of characters excluding, in one embodiment, semicolon, comma, and white spaces. If there is a need to place these characters in the VALUE, standard encoding methods, such as URL's type "%XX" encoding, can be used.

15 The "expires" attribute specifies a data string that defines the valid life time of the corresponding cookie. Once the expiration date has been reached, the cookie will no longer be stored in the client system. Thus, the client system will no longer respond to Web servers with the cookie. Many coding schemes for designating time can be used. In a preferred embodiment, the "expires" attribute is formatted as: Wdy, DD-Mon-YY HH:MM:SS GMT In this format, "Wdy" designates the day of a week, "DD-Mon-YY" designates the day, month and year, and "HH:MM:SS GMT" designates the hour, minute and second, in GMT time zone. Note that the "expires" attribute lets a client know when it is safe to purge a cookie, however, the client is not required to delete the cookie. If an expires attribute is not provided by the server, then the cookie expires when the user's session ends. This can be implemented by storing the cookie only in volatile memory.

35 The "domain=DOMAIN_NAME" attribute defines a domain for which the cookie is valid. The domain attribute is usually set using the domain name of the sending Web server. Client systems examine the domain attribute when making later http requests. If the server that the client system is accessing falls within the defined DOMAIN_NAME, then the cookie may be sent to the server when making the http request. (The "path" must also be examined as will be explained later.) When a server system falls within the defined DOMAIN_NAME, this is referred to as a "tail match." Note that a domain name that defines a subset of a domain is deemed to match a larger enclosing domain. For example, the host names "anvil.acme.com" and "shipping-.crate.acme.com" fall within the "acme.com" domain.

40 Only hosts within the specified domain can set a cookie for a domain. The value of the "domain" attribute must have at least two periods in them to prevent accepting values of the form ".com" and ".edu". If no domain name is specified, then the default value of the "domain" attribute is the domain name of the server that generated the cookie header.

45 The "path" attribute is used to specify a subset of file system directories in a domain for which the cookie is valid. If a cookie has already passed "domain" matching, then the path name of the URL for a requested document is compared with the "path" attribute. If there is a match, the cookie is considered valid and is sent along with the http request. All the characters of the defined path must match, however there may be additional characters on the path name. Thus, further defined subdirectories will match a path to the parent director. For example, the path "/foo" would match "/foo/bar", "/foo/bar.html", and even "/foobar", but "/foo" will not match the path "/". Note that the path "/" is the most general path since it will match any path. If no path is specified when a cookie is created, then the default path will be the same path as the document that was sent with the header which contains the cookie.

The last element of the cookie definition is the optional label of "secure." If a cookie is marked "secure," then the cookie will only be retransmitted if there is a secure communication channel to the server system. In a preferred embodiment of the present invention, this means that the cookie will only be sent to HTTPS servers. (HTTP over SSL) If the "secure" attribute is not specified, a cookie is considered safe to be sent over unsecured channels.

The defined extension to the HTTP protocol allows multiple set-cookie headers to be issued in a single HTTP response. Each set-cookie header should follow the conventions of the above described format.

Client Behavior

As previously described, when a client receives a set-cookie command in a header, the client system stores the cookie in some type of storage. In order not to place too much burden on client systems, each client system is expected to be able to store only a limited number of cookies. In one embodiment, the storage requirements for the client systems are:

- (1) 300 total cookies;
- (2) 4 kilobytes per cookie; and
- (2) 20 cookies per server or domain (note that this rule treats completely specified hosts and domains which are different as separate entities, and each entity has a 20 cookies limitation). Servers should not expect clients to be able to exceed these limits. When the 300 total cookies or the 20 cookie per server limit is exceeded, clients may delete the least recently used cookie even if the cookie's "expires" time has not passed.

If a cookie is received that matches the "NAME", "domain" and "path" attributes of a previously received cookie, then the previously received cookie will be overwritten. Using this technique, it is possible for a server to delete a cookie previously sent to a client. Specifically, a server that wishes to delete a previous cookie sends a cookie having "expires" time which is in the past that matches the "NAME", "domain" and "path" attributes of cookie to be deleted. Since the new overwritten cookie contains a expires time that has passed, the cookie will be deleted by the client system. Note "NAME", "domain" and "path" attributes of the expired cookie must match exactly those of the valid cookie. Since a system must be within the domain that is specified in the domain attribute, it is difficult for any server other than the originator of a cookie to delete or change a cookie.

When a client system that implements the present invention wishes to send an http request to a particular Web server, the client system first examines its cookie list to see if the cookie list contains any matching cookies that need to be sent to the particular Web server. Specifically, before the client sends an http request to a Web server, the client compares the URL of the requested Web document against all of the stored cookies. If any of the cookies in the cookie list matches the requested URL then information containing the name/value pairs of the matching cookies will be sent along with the HTTP request. The format of the line is:

Cookie: NAME1=value1; NAME2=value2; . . .

When a client sends cookies to a server, all cookies with a more specific path mapping should be sent before cookies with less specific path mappings. For example, a cookie "name1=foo" with a path mapping of "/bar" should be sent before a cookie "name2=foo2" with a path mapping of "/" if they are both to be sent since the path "/bar" is more specific than the global matching path "/".

Paths having a higher-level value do not override more specific path mappings. If there are multiple matches for a

given cookie name, but with separate paths, all the matching cookies will be sent. Thus, both the cookie "name=foo" with a path mapping of "/bar" and the cookie "name=foo" with a path mapping of "/" should be sent since they have different path names.

Some clients access Web servers on the Internet through firewall systems that are designed to prevent unwanted Internet traffic from affecting a local area network coupled to the Internet. Firewall systems often implement "proxy servers" that filter traffic to and from the Internet. It is important that proxy servers not cache Set-cookie commands when caching HTTP information. Thus, if a proxy server receives a response that contains a Set-cookie header, the proxy server should immediately propagate the Set-cookie header to the client. Similarly, if a client system request contains a "Cookie: " header, the cookie header should be forwarded through a proxy even if a conditional "If-modified-since" request is being made.

To further describe the present invention, the following examples describe a set of Web transactions operating in accordance with the present invention:

EXAMPLE 1

A client system requests a Web document from the Web server "telemarketing.acme.com" and receives in response:

Set-Cookie: CUSTOMER=WILE_E_COYOTE; path=/; expires=Wednesday, 9-Nov.-1999 23:12:40

The client system stores this cookie in a local (client-side) storage unit (e.g. mass storage 127 or memory 125). Since no domain name was specifically identified, the domain will be set to "telemarketing.acme.com" since that is the domain name of the server that generated the cookie. When the client later makes an http request for a document in any path (since the path is "/") of a server system in the telemarketing.acme.com domain, the client sends:

Cookie: CUSTOMER=WILE_E_COYOTE

Assuming the client system makes another request to the telemarketing.acme.com domain, the client might receive another cookie from the server such as:

Set-Cookie: PART_NUMBER=ROCKET_LAUNCHER;

path=/

The client will locally store this additional cookie. Again, no domain name was identified, such that the default domain, "telemarketing.acme.com" will be stored. Now, if the client makes yet another request to the "telemarketing.acme.com" domain, the client will send all the cookies it has for that domain. Specifically, the client sends:

Cookie: CUSTOMER=WILE_E_COYOTE;

PART_NUMBER=ROCKET_LAUNCHER

Assuming, the client continues transactions with the "telemarketing.acme.com" server, it may receive the following cookie from the server:

Set-Cookie: SHIPPING=FEDEX; path=/foo Then, if the client requests a document in path "/" on the "telemarketing.acme.com" server, the client will send two cookies as state information:

Cookie: CUSTOMER=WILE_E_COYOTE;

PART_NUMBER=ROCKET_LAUNCHER

Note that the cookie SHIPPING=FEDEX was not sent because the path "/" does not match the path "/foo". On the other hand, when the client requests a document on the "telemarketing.acme.com" server in path "/foo" on this server, then the client will send three cookies as state information:

Cookie: CUSTOMER=WILE_E_COYOTE;
PART_NUMBER=ROCKET_LAUNCHER;
SHIPPING=FEDEX

EXAMPLE 2

Assume that all of the transactions of Example 1 have been cleared. A client system then requests a Web document from the Web server "telemarketing.acme.com" and receives in response:

Set Cookie: PART_NUMBER=ROCKET_LAUNCHER_1;

path=/

The client stores this cookie in a local (client-side) storage unit. Since no domain name was specifically identified, the domain will be set to "telemarketing.acme.com". When the client later makes a request to a document in any path of a system in the telemarketing.acme.com domain, the client sends back the following data as information:

Cookie: PART_NUMBER=ROCKET_LAUNCHER_1

Assuming the client continues to access the "telemarketing.acme.com" server, the client may later receive from the server:

Set-Cookie: PART_NUMBER=RIDING_ROCKET_23;

path=/ammo

The new cookie has the same name (PART_NUMBER) as an old cookie stored on the client system. Note that the old cookie is not overwritten since the new cookie has a different path attribute. Now, if the client makes a request for a document in the path "/ammo" on the "telemarketing.acme.com" server, the client should send the following two cookies as state information:

Cookie: PART_NUMBER=RIDING_ROCKET_23;
PART_NUMBER=ROCKET_LAUNCHER_1

Both cookies are sent since the path of the requested document (""/ammo") matches both the "/" path of the first cookie and the "/ammo" path of the second cookie. Note that the cookie PART_NUMBER=RIDING_ROCKET_23 is sent first since it has a more specific path (""/ammo") than the global path ("/") associated with the cookie PART_NUMBER=ROCKET_LAUNCHER_1.

An On-line Shopping System Application

To illustrate one possible use of the state information system of the present invention, an implementation of an on-line shopping system will be described. The on-line shopping system allows customers to shop in one or more stores that are implemented as Web servers on the Internet. A customer can browse information on the Web servers that describe products available from the stores. When a desired product is found, the user can place the product into a "virtual shopping basket." The virtual shopping basket is implemented as a set of cookies that are sent to the client computer system and stored on the client computer system. At check-out time, the customer pays for the selected products using some type of payment system such as a credit card. After payment is received, the on-line shopping system notifies the stores to ship the selected products to the customer.

FIG. 5 is a flow chart showing the operation of the merchant system during an on-line shopping "trip" by a customer. The customer can run a browser on a client computer system, such as computer system 140 shown in FIG. 6 or client system 120 shown in FIG. 1B. The computer

system 140 of FIG. 6 includes a display device 142 (such as a monitor), a display screen 144, a cabinet 146 (which encloses components typically found in a computer, such as CPU, RAM, ROM, video card, hard drive, sound card, serial ports, etc.), a keyboard 148, a mouse 150, and a modem 152. Mouse 150 have one or more buttons, such as buttons 154. The computer needs to have some type of communication device such as that Modem 152 allows computer system 140 to be connected to the Internet. Other possible communication devices include ethernet network cards.

The customer uses Web browser software to access an on-line "merchant" server that is operated by a merchant having products to sell. This merchant server is a server computer system such as server system 122 shown in FIG. 1B. Specifically, the browser software sends an http request for the home Web page of a merchant Web server (step 212). The merchant Web server responds to the request with an HTML document that is displayed by the browser (step 214). The home Web page contains information about the merchant and its products (e.g., shoes, hats, shirts, etc.). The home Web page can implement a set of linked Web pages that describe the products that are available from the merchant. Each product may be associated with its own HTML document that fully describes the product. Products can be described using text, images, sounds, video clips, and any other communication form supported by Web browsers. The user can continue browsing through Web pages of the merchant server by repeating steps 212, 214, and 215.

After browsing through the Web pages provided by the server, the customer may select a product (step 216) by, for example, "clicking" (in the conventional manner) on an image of a product that causes the browser to request a Web page that fully describes the product. If the customer wishes to buy shoes from the merchant, the customer could click on a "buy it" button. The merchant server then sends an HTML form document that requests the customer to send necessary details for the purchase (step 218). For example, the customer may select a quantity, a desired style, and size of the product as requested by the form document. The browser then sends a POST command under HTTP, which transmits the data entered into the form to the merchant server (step 222). The data on the submitted form (e.g., quantity, size, style, etc.) is analyzed by the server and the transaction is processed. The server then generates a synthetic page and sends it to the browser running on the client system. This synthetic page preferably contains a thank you note along with confirmation information. Cookies containing information describing the selected product are also sent at this time (step 224).

The browser software running on the client system stores the cookies describing the selected products within the client computer system (step 226). The stored cookies include an identification of the contents of a virtual shopping basket that contains the products selected by the consumer. In an embodiment of the present invention, the cookies are stored in a file located in a storage medium (such as a hard disk) of client computer system 140.

The time interval for storing the cookies that describe the selected products can be set to any desired length. In one embodiment of the present invention, the cookies are deleted when the customer exits from the browser. This can be accomplished by not setting the "expires" attribute of the product description cookies. In another embodiment of the present invention, the cookies are kept valid (prior to their expiration) even after the customer exits from the browser and turns off computer 140. This can be accomplished by setting the "expires" attribute of the product description cookies to a later date.

13

After selecting a product, the customer may do additional shopping (e.g., buy a hat) from the same store or other stores (step 228). In this case, steps 212, 214, 215, 216, 218, 222, 224 and 226 need to be performed for the additional products. Each selection of a product in step 222 will result in the transmission of a cookie from the server to the client, which cookie identifies the selected product. The customer may also exit from the merchant system at any time.

When the customer desires to buy the products, the customer accesses a link that identifies a "check-out" Web page. The check-out Web page causes the browser to send all the product description cookies (230). Thus, the check-out Web page empties out the virtual shopping basket. The merchant server generates a total bill for all the products in the virtual shopping basket. The server may then request billing information (e.g., credit card number) and shipping (e.g., address) information from the customer using a form. In a preferred embodiment the transaction of credit card information is transmitted using a secure medium. The transaction server then performs a real-time credit card authorization. Once the transaction is authorized, transaction server sends messages to individual merchants to fulfill the order (step 240).

Other functions could be added to the above described merchant system. For example, several persons could use the same browser for shopping. In this case, the browser identifies the person doing the shopping, and assigns product description cookies to the appropriate person. Thus, each person would have their own virtual shopping basket.

The invention has been described with reference to specific exemplary embodiments thereof and various modifications and changes may be made thereto without departing from the broad spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense; the invention is limited only by the following claims.

What is claimed is:

1. A method of transferring state information between an http server and an http client, said method comprising the steps of:

requesting a file on said http server from said http client;
transmitting said file from said http server to said http client;
transmitting a state object from said http server to said http client; and
storing said state object on said http client.

2. The method of transferring state information as claimed in claim 1 wherein said state object comprises a name attribute.

3. The method of transferring state information as claimed in claim 1 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server and wherein said state object is transmitted along with said file.

4. The method of transferring state information as claimed in claim 1 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is within said domain.

5. The method of transferring state information as claimed in claim 1 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

14

6. The method of transferring state information as claimed in claim 1 wherein said state object includes an expiration attribute defining a valid life time of said state object.

7. The method of transferring state information as claimed in claim 1 wherein said state object includes an attribute requesting transmission using a secure channel.

8. The method of transferring state information as claimed in claim 1 wherein said state object is encoded within a header associated with said file.

9. A computer readable medium on an http client containing executable program instructions for performing a method comprising:

requesting a file on a http server;
receiving said file from said http server;
receiving a state object which specifies state information from said http server;
storing said state object on said http client.

10. A computer readable medium on an http server containing executable program instructions for performing a method comprising:

receiving a request for a file on said http server from an http client;
transmitting said file from said http server to said http client;
transmitting a state object which specifies state information from said http server to said http client.

11. A network of computer systems comprising:

a client system having a client processor and a client computer readable medium coupled to said client processor, said client computer readable medium containing program instructions for receiving a state object which specifies state information and for storing said state object on said client computer readable medium; a server system having a server processor and a server computer readable medium coupled to said server processor, said server system coupled to said client system through a network medium, said server computer readable medium containing program instructions for transmitting a file from said server system to said client system and for transmitting said state object to said client system.

12. A network as in claim 11 wherein said server system comprises an http server and wherein said client system comprises an http client.

13. A network as in claim 12 wherein said network medium comprises a client modem and a server modem and an interconnection between said client modem and said server modem.

14. A computer system, said computer system comprising:

a processor;
a memory coupled to said processor;
a computer readable medium coupled to said processor, said computer readable medium containing executable program instructions for:
requesting a file on a server;
receiving said file from said server;
receiving a state object which specifies state information from said server; and
storing said state object in one of said memory and said computer readable medium.

15. A computer readable medium as in claim 9 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server.

15

16. A computer readable medium as in claim 9 wherein said state object includes an expiration attribute defining a valid life time of said state object.

17. A computer readable medium as in claim 9 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

18. A computer readable medium as in claim 9 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is within said domain.

19. A computer readable medium as in claim 10 wherein said state object is transmitted from said http client to a server when said http client makes predefined http requests to said server.

20. A computer readable medium as in claim 10 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said http client to a server only when said http client makes an http request to said server and said server is within said domain.

21. A computer readable medium as in claim 10 wherein said state object includes an expiration attribute defining a valid life time of said state object.

10

15

20

25

16

22. A computer readable medium as in claim 10 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said http client to said server only when said http client makes said http request for a document within said path at said server.

23. A computer system as in claim 14 wherein said state object is transmitted from said computer system to a server when said computer system makes predefined requests to said server.

24. A computer system as in claim 14 wherein said state object includes a domain attribute defining a domain and said state object is transmitted from said computer system to a server only when said computer system makes a request to said server and said server is within said domain.

25. A computer system as in claim 14 wherein said state object further includes a path attribute defining a file system path and said state object is transmitted from said computer system to said server only when said computer system makes a request for a document within said path at said server.

26. A computer system as in claim 14 wherein said state object includes an expiration attribute defining a valid life time of said state object.

* * * * *

Computer Networking

A Top-Down Approach Featuring the Internet

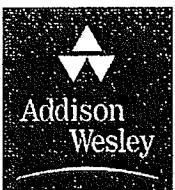
James F. Kurose

University of Massachusetts, Amherst



Keith W. Ross

Institut Érable



Boston San Francisco New York
London Toronto Sydney Tokyo Singapore Madrid
Mexico City Munich Paris Cape Town Hong Kong Montreal

Senior Acquisitions Editor Susan Hartman
Assistant Editor Lisa Kalner
Production Supervisor Patty Mahtani
Art Editor Helen Reebenacker
Executive Marketing Manager Michael Hirsch
Composition Pre-Press Company, Inc.
Technical Art PD & PS
Copyeditor Roberta Lewis
Proofreader Holly McLean Aldis
Cover Design Joyce Cosentino
Interior Design Delgado Design
Design Manager Regina Hagen
Cover Image © 1999 PhotoDisc, Inc.

Access the latest information about Addison-Wesley titles from our World Wide Web site:
<http://www.awl.com/cs>

The programs and applications presented in this book have been included for their instructional value. They have been tested with care, but are not guaranteed for any particular purpose. The publisher does not offer any warranties or representations, nor does it accept any liabilities with respect to the programs or applications.

Library of Congress Cataloging-in-Publication Data
Ross, Keith W., 1956-

Computer networking: a top-down approach featuring the Internet / Keith W. Ross,
James F. Kurose.

p. cm.

Includes bibliographic references and index

ISBN 0-201-47711-4

1. Internet (Computer network) I. Kurose, James F.
TK5105.875.I57 R689 2001

004.6—dc21

00-025295

Copyright © 2001 by Addison Wesley Longman, Inc.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior written permission of the publisher. Printed in the United States of America.

1 2 3 4 5 6 7 8 9 10-MA-020100

Error detection in the link layer is usually more sophisticated and implemented in hardware.

♦ **Error correction.** Error correction is similar to error detection, except that a receiver cannot only detect whether errors have been introduced in the frame but can also determine exactly where in the frame the errors have occurred (and hence correct these errors). Some protocols (such as ATM) provide link-layer error correction for the packet header rather than for the entire packet. We cover error detection and correction in Section 5.2.

♦ **Half-duplex and full-duplex.** With full-duplex transmission, the nodes at both ends of a link may transmit packets at the same time. With half-duplex transmission, a node cannot both transmit and receive at the same time.

As noted above, many of the services provided by the link layer have strong parallels with services provided at the transport layer. For example, both the link layer and the transport layer can provide reliable delivery. Although the mechanisms used to provide reliable delivery in the two layers are similar (see Section 3.4), the two reliable delivery services are not the same. A transport protocol provides reliable delivery between two processes on an end-to-end basis; a reliable link-layer protocol provides the reliable-delivery service between two nodes connected by a single link. Similarly, both link-layer and transport-layer protocols can provide flow control and error detection; again, flow control in a transport-layer protocol is provided on an end-to-end basis, whereas it is provided in a link-layer protocol on a node-to-adjacent-node basis.

5.1.2 Adapters Communicating

For a given communication link, the link-layer protocol is, for the most part, implemented in an adapter. An adapter is a board (or a PCMCIA card) that typically contains RAM, DSP chips, a host bus interface, and a link interface. Adapters are also commonly known as network interface cards or NICs. As shown in Figure 5.2, the network layer in the transmitting node (that is, a host or router) passes a network-layer datagram to the adapter that handles the sending side of the communication link. The adapter encapsulates the datagram in a frame and then transmits the frame into the communication link. At the other side, the receiving adapter receives the entire frame, extracts the network-layer datagram, and passes it to the network layer. If the link-layer protocol provides error detection, then it is the sending adapter that sets the error detection bits and it is the receiving adapter that performs error checking. If the link-layer protocol provides reliable delivery, then the mechanisms for reliable delivery (for example, sequence numbers, timers, and acknowledgments) are entirely implemented in the adapters. If the link-layer protocol provides random access (see Section 5.3), then the random access protocol is entirely implemented in the adapters.

An adapter is a semi-autonomous unit. For example, an adapter can receive a frame, determine if a frame is in error and discard the frame without notifying its

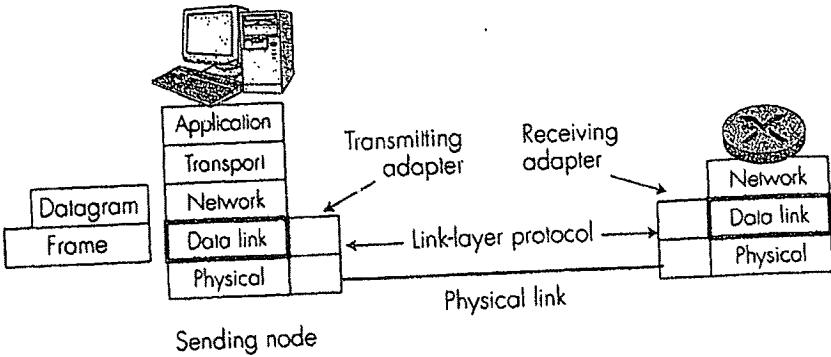


Figure 5.2 • The link-layer protocol for a communication link is implemented in the adapters at the two ends of the link

"parent" node. An adapter that receives a frame only interrupts its parent node when it wants to pass a network-layer datagram up the protocol stack. Similarly, when a node passes a datagram down the protocol stack to an adapter, the node fully delegates to the adapter the task of transmitting the datagram across that link. On the other hand, an adapter is not a completely autonomous unit. Although we have shown the adapter as a separate "box" in Figure 5.3, the adapter is typically housed in the same physical box as the rest of the node, shares power and busses with the rest of the node, and is ultimately under the control of the node.

As shown in Figure 5.3, the main components of an adapter are the bus interface and the link interface. The bus interface is responsible for communicating with the adapter's parent node. It transfers data and control information between the node and the NIC. The link interface is responsible for implementing the link-layer protocol. In addition to framing and de-framing datagrams, it may provide error detection, random access, and other link-layer functions. It also includes the transmit and receive circuitry. For popular link-layer technologies, such as Ethernet, the link interface is implemented by chip set that can be bought on the commodity market. For this reason, Ethernet adapters are incredibly cheap—often less than \$30 for 10 Mbps and 100 Mbps transmission rates.

Adapter design has become very sophisticated over the years. One of the critical issues in adapter performance has always been whether the adapter can move data in and out of a node at the full line speed, that is, at the transmission rate of the link. You can learn more about adapter architecture for 10 Mbps Ethernet, 100 Mbps Ethernet, and 155 Mbps ATM by visiting the 3Com adapter page [3Com 1999]. *Data Communications* magazine provides a nice introduction to Gbps Ethernet adapters [GigaAdapter 1997].

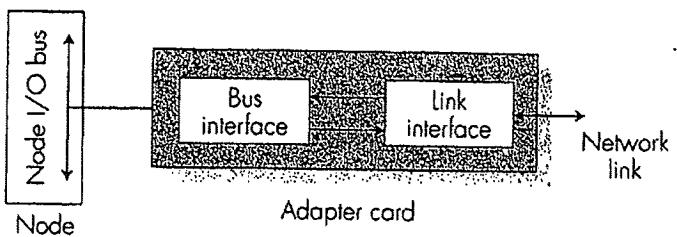


Figure 5.3 ♦ The adapter is a semi-autonomous unit

5.2 ♦ Error Detection and Correction Techniques

In the previous section, we noted that bit-level error detection and correction—detecting and correcting the corruption of bits in a data-link-layer frame sent from one node to another physically connected neighboring node—are two services often provided by the data-link layer. We saw in Chapter 3 that error detection and correction services are also often offered at the transport layer as well. In this section, we'll examine a few of the simplest techniques that can be used to detect and, in some cases, correct such bit errors. A full treatment of the theory and implementation of this topic is itself the topic of many textbooks (for example, [Schwartz 1980]), and our treatment here is necessarily brief. Our goal here is to develop an intuitive feel for the capabilities that error detection and correction techniques provide, and to see how a few simple techniques work and are used in practice in the data link layer.

Figure 5.4 illustrates the setting for our study. At the sending node, data, D , to be protected against bit errors is augmented with error-detection and correction bits, EDC . Typically, the data to be protected includes not only the datagram passed down from the network layer for transmission across the link, but also link-level addressing information, sequence numbers, and other fields in the data-link frame header. Both D and EDC are sent to the receiving node in a link-level frame. At the receiving node, a sequence of bits, D' and EDC' are received. Note that D' and EDC' may differ from the original D and EDC as a result of in-transit bit flips.

The receiver's challenge is to determine whether or not D' is the same as the original D , given that it has only received D' and EDC' . The exact wording of the receiver's decision in Figure 5.4 (we ask whether an error is detected, not whether an error has occurred!) is important. Error-detection and correction techniques allow the receiver to sometimes, *but not always*, detect that bit errors have occurred. That is, even with the use of error-detection bits there will still be a possibility that

Taking the Byte Out of Cookies: Privacy, Consent, and the Web

Daniel Lin

*Department of Computer Science
University of Illinois at Urbana-Champaign*

Michael C. Loui

*Department of Electrical and Computer Engineering,
Coordinated Science Laboratory, and Graduate College
University of Illinois at Urbana-Champaign*

Abstract: We consider the privacy of personal information on the World Wide Web, emphasizing a concept of privacy as an aspect of social relationships between individuals. We make three contributions to understanding the right to privacy on the Web: (1) we highlight the role of informed consent as an important consideration for privacy, (2) we identify conditions under which the collection and centralization of personal information can be ethically justified, and (3) we offer an interpretation of a "reasonable expectation of privacy" for Internet cookies, a mechanism used by Web sites to remember information about visits to that site.

The views, opinions, and conclusions of this paper are not necessarily those of the University of Illinois. Preliminary versions of this paper were presented at the Seventh Annual Meeting of the Association for Practical and Professional Ethics, Dallas, TX, February 26[28, 1998, and the ACM Policy 98 Conference, Washington D.C., May 10-12, 1998.

Address for correspondence: Michael C. Loui, Graduate College, 801 S. Wright Street, Champaign, IL 61820-6210, e-mail: m-loui@uiuc.edu, telephone: (217) 333-6715, fax: (217) 333-8019.

1 Introduction

What is our right to privacy on the World Wide Web? Is it reasonable for us to assume that we should always have control over our personal information on the Web? These questions become increasingly urgent as we continue to unknowingly release our personal information into cyberspace.

In a recent Internet privacy survey directed by Alan Westin, 53 percent of Internet users and 57 percent of online service users were concerned that their Internet browsing behavior would be linked to their e-mail addresses and disclosed to other people or organizations [1]. Such concerns about privacy are not new. Many of us express similar concerns when we apply for a credit card or a car loan. What is new about our privacy concerns, however, is the environment in which we express them — the Web. According to Moor [2], information on the Web is "greased" and its manipulation occurs more easily and at a much grander scale. Because the Web is a new environment, there are few transactions on the Web whose impact is completely understood by the general public. Furthermore, unlike applying for a credit card or car loan, we may not even be aware that we are releasing personal information as we "surf" the Web.

For instance, many Web sites use a mechanism called a "cookie" which silently collects information about our visit to that site. The consequence of dealing with a new environment like the Web is that we are unable to make deliberate decisions about revealing our personal information. Simply put, we cannot make well-informed decisions because we do not have reasonable expectations of privacy on the Web.

In this paper, we make three contributions to understanding privacy and the right to privacy on the Web. First, we highlight the role of informed consent in the theory of privacy. Because we currently do not have a reasonable expectation of privacy on the Web, informed consent is important: if we do not have enough knowledge to make an informed decision about revealing our personal information on the Web, we should be given that knowledge before making our decisions. Second, we establish the conditions under which the collection and centralization of personal information are ethically justifiable. The collection and centralization of personal information on the Web affect our privacy in fundamentally different ways. While previous notions of privacy describe when such manipulations of information cause losses or violations of privacy, they do not address how such violations might be avoided. We address this issue by showing

that informed consent is sufficient for the collection of personal information to be ethical. Third, we offer an interpretation of what a "reasonable expectation of privacy" may mean for Internet cookies. In doing so, we distinguish between morally permissible and immoral uses of cookies.

Section 2 discusses dominant approaches to privacy and the importance of privacy. Section 3 outlines the argument for informed consent. In Sections 4 and 5, we develop the ethical boundaries for collecting and centralizing personal information. Section 6 explores the notion of a reasonable expectation of privacy in public places. In Section 7, we introduce the idea that our reasonable expectations of privacy on the Web should reflect the realization that the Internet is a public place. Finally, in Section 8 we apply our moral analysis to the current privacy concern with Internet cookies.

2 Ethical Theories of Privacy

While at first glance, the concept of privacy may seem simple, many have struggled to adequately describe what privacy actually is. In this section, we present a critical analysis of different theories of privacy.

2.1 Privacy as Nonintrusion

In their 1890 Harvard Law Review article, "The Right to Privacy," Warren and Brandeis [3] wrote that privacy is the right "to be let alone." For example, Bob loses some privacy when Alice rummages through his desk. In such a scenario, Alice is not leaving Bob alone. By defining privacy as the right "to be let alone," however, we may relate privacy to situations which have no true connections to it. If Alice clubs Bob on the head with a baseball bat, she has not invaded his privacy. She has assaulted him. Nevertheless, she has indeed violated his right "to be let alone."

2.2 Privacy as Control of Information

Fried [4], Westin [5], and Beardsley [6] define privacy as the control of personal information. That is, if we can determine how much personal information we can reveal and to whom we reveal that information, we can prevent violations of our privacy. If Alice rummages through Bob's desk, finds his credit records, and reveals this information to Charles, then Bob has lost control of his credit information and suffers a violation of privacy. However, while control is an important aspect of privacy, privacy solely as control may erroneously describe instances which do not involve privacy. That is, there are many situations in which people lose or never have control of their personal information yet suffer no violation of privacy. Moor [9] uses the example of medical records. Bob's medical records may be passed on to various doctors and nurses so that he receives proper medical care. While Bob has no control over this medical information, he does not suffer a violation of privacy.

2.3 Privacy as Undocumented Personal Knowledge

Seeking a better definition of privacy, Parent [7] defines privacy as the condition in which undocumented personal information is not possessed by others. For Parent, personal information consists of those facts which people do not wish to reveal about themselves. It also includes facts about which a person may be very sensitive about, such as weight or height. Undocumented information is any information which cannot be found in public documents such as newspapers or court proceedings. From these definitions, it follows that any personal information which is documented must have once been undocumented. Parent acknowledges that when personal information is first published, there is an invasion of privacy. He calls later uses of such information "gratuitous exploitation" rather than further violations of privacy. Suppose Alice is sunbathing naked on her private beach. If Bob takes a photograph of Alice and Star magazine prints the photograph, this personal information now becomes documented information. According to Parent, knowledge of Alice's nude body is now publicly available, and therefore no privacy is lost the next time someone sees Alice nude. This conclusion seems counterintuitive to us.

2.4 Privacy as Restricted Access

Perhaps the most complete conception of privacy is based on restricted access, due to Gavison [8]. She introduces three aspects of privacy: Secrecy: The extent to which we are known to others. Anonymity: The extent to which we are the subject of others' attention. Solitude: The extent to which others have physical access to us. A loss of privacy occurs when the degree of our secrecy, anonymity, or solitude decreases. Gavison makes an important distinction between the loss of privacy and the violation of privacy. That is, losses of privacy are not necessarily undesirable. Each situation must be assessed to determine whether the loss of privacy limits the functions of privacy. For Gavison, such functions of privacy include the following:

1. The creation of an environment where trust, love, friendship, and intimacy can be maintained.
2. Freedom from physical access.
3. Promotion of liberty of actions. By this we mean that privacy permits individuals "to do what they would not do without it for fear of an unpleasant or hostile reaction from others." Specifically, it promotes our mental health and autonomy and protects us from censure and ridicule.

Moor [9] also adopts the framework of privacy as restricted access. He describes two kinds of privacy: natural privacy and normative privacy. Loss of natural privacy is not necessarily an invasion of privacy. If Bob is meditating in the Grand Canyon, he is enjoying his state of natural privacy. If a group of noisy tourists riding mules descends upon him, however, he cannot claim an invasion of privacy. On the other hand, Bob may be relaxing in his New York apartment, smoking a pipe and reading the paper. If the same

group of noisy tourists riding mules peers into his window, they are violating his normative right to privacy.¹

2.5 The Value of Privacy

These different theories agree that privacy is an important part of our lives. The value of privacy has been described both as a means to achieve or maintain other important goals of our lives and as an end, having intrinsic or inherent value, itself. Rachels [10] considers privacy as a means. For him, privacy is important because it enables us to create social context in our relationships with other people:

The value of privacy is based on the idea that there is close connection between our ability to control who has access to us and who has information about us, and our ability to create and maintain different sorts of relationships with different people. For instance, Bob may be aggressive, relentless, and ungenerous in dealing with his business partners. At the same time, he may be tender and sensitive with his family and friends. Privacy, as a means, allows us to exhibit different patterns of behavior with different people, to maintain different social relationships. Benn [11] argues that privacy has intrinsic value. He claims that there is a general principle of privacy based on respect for persons. He begins by stating that people have a general liberty to do what they choose unless someone else has good reason for preventing it.²

For Benn, the principle of privacy offers such a reason to limit the general liberty of others to observe and report at will. For instance, imagine that Bob is unhappy that Alice is watching him. Bob's behavior is fundamentally changed by the mere fact that he knows Alice is watching him. Such unwanted observation does not treat Bob with the respect he deserves. Suppose, on the other hand, that Bob did not know that Alice was watching him. He would behave as if he were being unwatched, which is not the way he would want to behave. Covert observation is objectionable, as Benn states, because it deceives Bob about the context of his actions and treats him without dignity.

3 Ethical Theories of Informed Consent

Many violations of privacy could be avoided if an element of consent or awareness was introduced. Consider again, for example, the scenario presented in Section 2.3 about the theory of privacy as undocumented knowledge. The main problem, it seems, is that Bob did not ask Alice if she would mind being photographed. Similarly, Alice did not ask for permission before she rummaged through Bob's desk in Section 2.1. In Section 2.2, Alice did not think about Bob's feelings before she blurted out his credit record to all her friends. In each of these scenarios, had the victim been properly informed and given a choice, a violation of privacy might not have occurred. Consideration for the victim is the essence of the theory of informed consent.

To our knowledge, previous accounts of privacy have not emphasized the role of consent. Instead, they attempt to define precise boundaries around privacy. These accounts explain privacy by defining the circumstances under which an individual's privacy has been compromised. That is, they isolate the concept of privacy solely as an issue of safekeeping information, without consideration of social context. For example, Parent's [7] conception of privacy as undocumented knowledge precisely delineates what is private (undocumented personal knowledge) and what is not (documented personal knowledge). By introducing the concept of consent, our approach in this paper is to emphasize the concept of privacy as an aspect of the social relationship between individuals. Had Alice consented to being photographed after being fully informed about potential consequences, we do not believe she would have still suffered an invasion of privacy. Our approach essentially builds upon Gavison's [8] neutral concept of privacy. Recall her argument that different situations must be examined to determine whether a loss of privacy results in a violation. We will argue in Section 4 that linking consent to privacy provides a framework for deciding whether a loss of privacy may be a violation of privacy. In the following sections, we present a brief summary of the theory of informed consent in both medicine and engineering.

3.1 A Brief History of Informed Consent

The theory of informed consent has been most developed in the medical ethics. Historically, physicians did not believe that the consent of the patient was necessary. Consider Hippocrates's advice to future physicians. Perform [these duties] calmly and adroitly, concealing most things from the patient while you are attending him. Give necessary orders with cheerfulness and serenity, turning his attention away from what is being done to him; sometimes reprove sharply and emphatically, and sometimes comfort with solicitude and attention, revealing nothing of the patient's future or present condition. The term "informed consent" first appeared in 1957 in the decision *Salgo v. Leland Stanford, Jr. University Board of Trustees*. In that case the court asked whether the patient had been adequately informed before consenting to the medical procedure. By the 1970's and 1980's, concerns with civil rights, women's rights, the consumer movement, and the rights of prisoners and the mentally ill had brought informed consent to the attention of the medical community.

3.2 Informed Consent in Medical Ethics

Faden and Beauchamp [12] present five basic elements necessary for informed consent in medicine. Disclosure: All pertinent information must be disclosed to the patient before he makes a decision. This includes information which the patient may not have requested. Comprehension: The patient must understand the general nature of the illness, the risks and benefits of each treatment, and the reasons for and

against these options. Voluntariness: The patient should be under no pressure or duress when making the decision. Competence: The patient should be able to take responsibility for making the decision (not a child or someone who is severely mentally ill). Consent: The patient must be given the choice to decide which alternative to take. Gorovitz [13] feels that informed consent is necessary: If the patient understands what the physician proposes to do, and thus informed, consents to its being done, then the medical intervention is not imposed on the patient in violation of the patient's autonomy; rather, that medical intervention is properly viewed as a service provided to the patient at the patient's request. By providing the patient with a choice in his own treatment, informed consent promotes individual autonomy and encourages rational decision making. It also protects the physician against charges of assault from a patient who did not want a particular treatment. Like the theory of privacy, the theory of informed consent is formally grounded on the principle of respect for persons.

3.2.1 Against Informed Consent

Some physicians believe that informed consent is not only impractical, but may even be detrimental to the patient. They argue that it is not possible for the patient to develop a sufficient comprehension of the illness as dictated by the doctrine of informed consent. After all, physicians go through years of medical school and residency to develop such an understanding. Further, physicians have neither time, skills, nor sensitivities to properly inform a patient. Additionally, studies have shown that patients, even when properly informed, often do not remember what they have been told. Most likely, these physicians argue, the patient did not want to be given so much information in the first place. It is also argued that obtaining informed consent may even be dangerous to the patient. With limited understanding, patients may make decisions which are detrimental to their own well-being. They may even develop unnecessary fears and anxieties from a more thorough understanding of the potential risks and discomforts of the treatments. None of these arguments seem strong enough to counter the principle of respect for persons. In fact, a wrong "medical" decision may not be the wrong decision from a broader perspective of the patient's life. For example, Bob may decide to forgo an operation on his larynx because he makes his living as an opera singer. While his decision may be medically unwise, Bob knows that such an operation would ruin his career and even his life. As Gorovitz [13] writes, "the right to choose is not limited to the right to choose rightly."

3.2.2 Exceptions to Informed Consent

Lidz [14] identifies situations in which it may be acceptable not to obtain informed consent: Emergency: There is not enough time to obtain informed consent without seriously risking the well-being of the patient. Incompetence: The pa-

tient is unable to make a decision for the situation at hand. The patient may be, for instance, intoxicated, unconscious, or senile. In such a situation, it is appropriate to secure a surrogate who may make decisions for the patient. Waiver: Under no pressure to do so, the patient waives the right to informed consent. Therapeutic Privilege: This exception allows the withholding of information that the physician feels would be "harmful." Therapeutic privilege prevents violation of the physician's primary duty of doing what is beneficial for the patient solely because of a legal duty to obtain informed consent. Unfortunately, many different interpretations of this privilege exist, ranging from the most stringent to the most lenient.

3.3 Informed Consent in Engineering Ethics

Martin and Schinzingher [15] bring the theory of informed consent into the area of engineering ethics. They believe that socially responsible engineers must consider the informed consent of those who will use their products. Engineers should realize the consequences of designing and manufacturing a product. Customers must be given all pertinent information about a product so that they can rationally decide whether to purchase it. For instance, if Bob creates a new security device which identifies employees by scanning their retinas with a laser, he must reveal all the risks as well as benefits to those who are interested in the device. Martin and Schinzingher emphasize that information must be voluntarily presented even if the customer has not requested it. Therefore, even if the customer is not interested, Bob, as a responsible engineer, would be obligated to reveal the risks of his product.

4 Collection of Information

Computers have entirely changed the way we collect information. According to Johnson [16], not only has computer technology increased the scale of information gathering, it has also enabled new kinds of information to be collected. Without computer technology, for instance, the scientific information collected from NASA's recent Mars mission would not have been possible. In today's modern environment, we would likely find it impossible to eliminate or even limit the collection of personal information. Rather, our efforts should concentrate on balancing the privacy rights of the individual with the needs of the community to collect information. In this section, we link previously discussed theories of privacy and consent to provide a framework for determining when the act of collecting personal information is ethical.

4.1 Distinguishing Collection from Use

It is important to distinguish the collection of personal information from its use. When we do not distinguish collection from use, we may easily blame the tool which collects information for undesired consequences. Mainstream mov-

ies such as The Net and Hackers have popularized the notion that computers and the Internet are the root of privacy and information problems. Numerous individuals have been arrested or denied service because databases contained incorrect or inaccurate information. For example, according to Forester and Morrison [17], Barbara Ward was continuously refused accommodation by landlords because her name was in a database of names of people who had taken their landlords to court.³ Quittner [18] mentions examples of computer misuse, including Sara Lee's plan to collaborate with Lovelace Health Systems to match employee health records with work performance reports to find workers who might benefit from antidepressants. These examples, however, do not reveal anything about the actual ethics of the collection of information. They are, rather, consequences of the use or misuse of collected information.⁴ While it is important to examine ethical decisions in the use of information, we wish to explore the ethics of the act of collecting personal information.⁵

4.2 When Is Collecting Personal Information Ethical?

When does the collection of information result in the loss of privacy? Under Gavison's definition of privacy as restricted access, the collection of information results in a loss of privacy when our degree of secrecy or anonymity is compromised. We might term the collected information which results in such a loss of privacy as personal information. Nevertheless, such a loss of privacy is not necessarily a violation of privacy. Recall Benn's assertion [11] that privacy can be justified under the respect for persons argument. Therefore as long as the principle of respect for persons is maintained, the collection of personal information can indeed be ethical. Consider the question raised by Benn [11]: How reasonable is it, then, for a person to resent being treated much in the same way that a birdwatcher might treat a redstart? It is reasonable for Bob to resent being watched like a bird because he has no knowledge that he is being watched. Even if Bob knew he was being watched, he probably would not like to be treated like an animal or specimen. Likewise, Bob probably would not be happy if someone read his private journal (a collection of personal information) without his permission. In all these examples, poor Bob is being treated neither with proper respect nor with regard for his dignity. One way to collect personal information from Bob while still treating him with proper respect is to obtain his informed consent. Section 3.2 discussed the basic elements needed for informed consent in medical ethics. When we extend informed consent from medicine to the collection of personal information, the elements of disclosure and comprehension must include the consequences of revealing or not revealing personal information: how the information will be used, who will have access to the information, how long the information will be kept, etc. Only with this information can a truly informed choice be made. If Bob understands the

reasons we would like to obtain his personal information and knows the consequences of revealing such information, then he can make an informed choice whether to divulge such information. Obtaining Bob's informed consent shows that we respect him as an autonomous being capable of making rational decisions. Both the theory of informed consent and the theory of privacy can be grounded on the principle of respect for persons. Since informed consent preserves this principle of respect for persons, it ensures that a collection personal information will not result in a violation of privacy. 4.3 When Is Collecting Personal Information Unethical? Does collecting personal information without obtaining consent necessarily result in a violation of privacy? That is to say, while obtaining informed consent satisfies the principle of respect for persons, a lack of informed consent may not necessarily show a disregard for it. Recall Benn's [11] assertion that only a general principle of privacy can be based on the principle of respect for persons: General principles do not determine solutions to moral problems of this kind. They indicate what needs to be justified, where the onus of justification lies, and what can count as justification [11]. Such a general principle offers only a minimal right to immunity.⁶

We believe that the concept of a reasonable expectation of privacy answers Benn's questions of what needs to be justified, where the onus of justification lies, and what can count as justification. For instance, in order to show that Alice should not be able to observe Bob on grounds of general liberty, Bob must have a reasonable expectation of privacy not to be observed in that situation. If Bob does not have a reasonable expectation of privacy, then Alice does not need to justify her desire to observe Bob, and therefore obtaining informed consent does not seem necessary. If, however, Bob does have a reasonable expectation of privacy, and Alice still insists on observing him, then a violation of privacy does occur. In this case, Alice has collected personal information in an unethical manner. Informed consent can transform a potentially unethical collection of personal information into an ethical one. A collection of personal information is unethical when it does not comport with the reasonable expectation of privacy for the situation at hand. Obtaining informed consent in these situations is necessary to avoid violations of privacy. Not all collections of personal information exceed the boundaries of a reasonable expectation of privacy, however. That is, obtaining informed consent may be sufficient for an ethical collection of data, but it is certainly not necessary. Alice does not need Bob's informed consent to watch him walk through a public square because Bob's reasonable expectation of privacy in this situation is limited. He would not reasonably expect that his right to privacy would protect him from observers as he walks through a public square. Suppose, however, that Alice peers into Bob's apartment window. Such an action obviously does not fall within Bob's reasonable expectation of privacy. Alice is en-

gaging in an unethical collection of personal information, and therefore violating Bob's privacy. Had Alice requested permission to observe Bob, however, a compromise may have been reached.⁷ Thus, Bob may risk the loss of privacy, if Alice seeks prior consultation before her attempt to observe him. In such a situation, the principle for respect for persons is preserved, Alice is free to conduct her observations, and Bob does not suffer a violation of privacy.

5 Centralization of Information

In Section 4, we established that the collection of personal information does not necessarily result in an unethical loss of privacy. Specifically, if Bob gives his informed consent before releasing any personal information, the loss of privacy which he experiences is not a violation. When Bob does give his informed consent, he essentially weighs all the perceived costs and benefits and concludes that the desirable functions of privacy (environment for maintaining relationships, freedom from physical access, and liberty of actions) are still preserved. In this section, we distinguish the act of collecting information from the act of centralizing information. By centralization of information, we mean the process of aggregating large quantities of personal information which have been collected for different purposes, and using this aggregation of information for an entirely new purpose. We will argue that the centralization of information, in contrast to its collection, is unethical because centralization contravenes the desirable functions of privacy identified in Section 2.4.

5.1 Creation of a Dossier Society

In 1986, Laudon [20] warned about the danger of the dossier society, in which personal files from different government agencies are integrated into a permanent national database. Advances in technology in addition to the government's loose interpretation and enforcement of the protections of the Privacy Act of 1974 had made the dossier society a real possibility.⁸ Not only did the development of a dossier society threaten the traditional American desire for a limited role in government, it had harmful cultural consequences as well:

From the individual's point of view, the most significant characteristic of the dossier society is that decisions made about us as citizens, employees, consumers, debtors, and supplicants rely less and less on personal face-to-face contact, on what we say, or even on what we do. Instead decisions are based on information that is held in national systems, and interpreted by bureaucrats and clerical workers in distant locations. The decisions made about us are based on a comprehensive "data image" drawn from diverse files [20].

What concerned Laudon most was the aggregation of power that a dossier society would bring to the federal gov-

ernment. He feared the necessity of explaining an "official life" to any government official who demanded such an explanation. Such concerns led him to criticize the potential development of the FBI's plan to create a national computerized criminal history into a general purpose national information database. Twelve years later however, we realize that fears about the dossier society are not limited solely to "official" information held by the federal government. The integration of commercial databases in industry has also created a de facto national database which is less official and more personal, even detailing preferences and habits. Culnan and Smith [22] describe the public uproar in 1991 over Lotus Marketplace: Households, a CD-ROM database which contained information, including lifestyle and purchasing propensities, of 120 million individuals in 80 million U.S. households. In June 1996, Lexis-Nexis launched its P-Trak service, advertised as the digital equivalent of the phone directory [23]. The P-Trak service, again, caused public concern because it gave access to Social Security Numbers. Within days of its launch, Lexis-Nexis, realizing the potential dangers of revealing such personal information, removed access to the Social Security Numbers.

5.2 Economic Theory of Information

The Lotus Marketplace: Households and Lexis-Nexis P-Trak controversies showed that centralized information could be easily accessible to everyone. Economic theory suggests that the cost of acquiring information guides behavior. For instance, if the perceived cost to Bob of learning how to select the very best apple in the supermarket is greater than its perceived benefits, Bob may decide to select a relatively good apple rather than the very best. Decreasing the cost of acquiring information, easily accessible databases increase the chance that persons will search for information that they would not otherwise seek because the cost would have been too high. Because there is such a low cost for obtaining the information, people may even acquire information which is neither pertinent nor reliable for the decisions they need to make. Imagine if Bob's business competitors and his intimate friends could easily obtain the same comprehensive data image which detailed not only his tax and credit records but his culinary preferences and purchasing habits as well. Easy access to an integrated federal and industry database seems devastating to our notions of privacy.

5.3 Violation of Privacy Without Loss of Privacy

Is there an ethical basis to our fear of the centralization of personal information? To the extent that no extra information is collected in the centralization process, how can there be any additional loss of privacy? While there may not necessarily be a loss of privacy, there is an unethical violation of privacy. Samet's [24] points out that we consider it a violation of privacy if Bob looks into our window and takes note of what we are doing. However, we experience no violation

of privacy if Bob looks out of his own window and notices what we are doing outside.⁹

In addition, let's assume all of Bob's family and friends also record what they see about us out their windows. Later that week, they all get together to share and compare notes. There is something disturbing about the detailed personal profile that Bob's family and friends would be able to compile.

According to Rachels [10], privacy is important because it provides the proper context for us to create and maintain different sorts of relationships with different people. Centralization of personal information is unethical because it destroys this important context which privacy provides. The comprehensive data image created in a centralized database denies us the ability to engage in the "different patterns of behavior associated with different relationships." While we have not consented to divulging certain information to certain parties, all parties who use the comprehensive data image can acquire that information. We might even say that a comprehensive data image is a portrayal of person that really does not exist. That is to say, at no single moment in time, do we display all behaviors that may be revealed in the comprehensive data image. Rather, we maintain different faces and behaviors in different contexts and relationships. Further, such a comprehensive data image permits subsets of information that are not personally or socially meaningful. The centralization of information creates an unsound profile because it takes information out of its original context and uses it in a different context. While certain information may be accurate in the context for which it was initially collected, it may be inaccurate when it is moved into a different context in a process of centralization. We may observe Bob as a careless free spirit at the racy nightclub which he frequents on the weekends. This certainly does not mean that he exhibits the same characteristics in his decisions as CEO of his company or as a father to his children. For these reasons, we conclude that centralization violates privacy. By changing the context in which information was initially collected, it counteracts a key function of privacy: the ability to create and maintain different relationships.¹⁰ 5.4 Informed Consent and Centralization of Information In section 4.3, we suggested that obtaining informed consent is sufficient for an ethical collection of information. How does obtaining informed consent affect the centralization of information? As discussed in the previous section, we believe that the centralization of information is *prima facie* unethical. Unlike the collection of information, which can be ethical even without obtaining informed consent, there are no conditions under which centralization of information can be ethical without obtaining informed consent. In fact, in situations where personal information is centralized, it is often impractical or impossible to even obtain the informed consent of the individuals.

For example, Parker [25] describes the ethical conflict faced by a scientist who finds two different kinds of data on the same subject pool. The scientist believes that there would be significant scientific value in merging and analyzing the data. Unfortunately, while informed consent had been obtained from the subject pool for collecting their personal information for the earlier studies, the subjects have now dispersed, and it is impossible to obtain their informed consent for this centralization of information. By merging the data, the scientist would be unethical. Parker suggests that the solution is to obtain the informed consent of a proxy, or representative of the subject pool, such as an independent committee. The proxy would weigh the benefits of the research against the violation of privacy to decide whether the merging should be allowed.

Parker's solution of a proxy may be sufficient for his scenario because the amount of centralization is small. We do not believe, however, that the solution of obtaining informed consent through a proxy can be extended to scenarios with large amounts of centralization. That is, when the scale of centralization is large enough to create a digital dossier, it is unlikely that an individual will consent to such a portrayal. Such profiles are dangerous because the superfluous information they provide affects how decisions are made. Furthermore, large scale centralization of information is a violation of privacy because the data images it creates are used out of context, and counteract the ability to create and maintain different relationships.

6 Reasonable Expectations of Privacy in Public Places

From the foregoing discussion, it does not seem that we necessarily have a *prima facie* right to privacy outside of that based on Benn's general principle of privacy. In fact, Ware [26] suggests that society will not and should not protect an individual's privacy at all costs. There is inevitable conflict between an individual's right to privacy and a community's rights. Such conflict may be manifested in monetary terms, inefficiencies to systems, or denial of desirable social services. This awareness of community rights is particularly pertinent to the right to privacy in public places, a concept introduced by Nissenbaum [27].

6.1 Privacy in Public Places

For Nissenbaum, current theories of privacy (Section 2) do not adequately address the relevance of privacy in realms other than the intimate and personal. Previous works on privacy ([4], [5], [6], [7], [28]) attempt to define a distinct and mutually exclusive boundary between an intimate personal realm where privacy is protected and a public realm where privacy has no relevance and all information is available to everyone. Nissenbaum [27] points out that while "there is a broad consensus on what information may be classified

personal and intimate, there is, on the other hand little, if anything, that people universally would admit into a completely public realm if by that we mean that it is governed by no norms of privacy whatever." Rather than viewing a context as purely private or purely public, we might view the context as having both private and public elements. Nissenbaum criticizes the notion that there is a relationship between a place and information which might be obtained in that place. That is, we should not assume that information is public simply because the place in which it was obtained is public. For instance, "shoppers may not object to using open shopping carts but may sense violation if inquisitive neighbors noted and reported on their purchases" [27]. For Nissenbaum, a conception of privacy should extend consideration to all information, including information which is obtained in public places.

6.2 Reasonable Expectation of Privacy

The protection of privacy in public realms cannot be as strong as in intimate realms. While we should extend consideration of privacy to the public realm, we must also acknowledge the limitations of privacy in the public realm. Moor's [9] concept of a naturally private situation captures the essence of the limitations of our right to privacy. A naturally private situation seems to be nothing more than the experience of isolation in a public place. Recall the sudden interruption of Bob's peaceful meditation in the Grand Canyon by the rude mule riding tourists. Although the intrusion is annoying and unfortunate, Bob has no right to privacy in this public place. The protection of privacy in such a public realm is limited, and his loss of privacy cannot be a violation.

How, then, do we determine what is reasonable in our expectation of privacy in public places? Bob may experience a loss of privacy when he purchases some books at the local bookstore because the bookstore may record Bob's purchasing preferences for inventory purposes. Such a collection of personal information would reduce Bob's secrecy and anonymity. The collection of this information seems reasonable because such public transactions are a necessary aspect of everyday life in society. Nevertheless, Bob may also reasonably expect that the information obtained by the bookstore will not be shared with third parties because such sharing was not the initial purpose for collecting the information. Our expectation of the protections of privacy should consider the needs of society's other entities such as industry, government, and community. A reasonable expectation is one which considers the practical harms which accompany a loss of privacy. That is, a claim of privacy seems reasonable only if there might be potential harm caused by the loss of privacy. For instance, Bob's claim of privacy would not be reasonable if Alice merely observed the nice red shirt he was wearing. A reasonable expectation of privacy should also consider the dynamic nature of the hierarchy of rights: depending upon the context of the situation, there may be other

intrinsic principles and rights which may be more important than privacy. When Bob is in his house, his right to privacy may take precedence over Sally's right to observe him. If he is in a public square, however, Sally's general liberty to observe may take precedence over Bob's right to privacy. Another example of the dynamic nature of rights is the revelation Bob's past history of drug addiction in his criminal trial. In normal circumstances, such a revelation may be inappropriate and a violation of Bob's privacy. In a criminal trial, however, the principle of social justice takes precedence over the principle of individual privacy. Reasonable expectations of privacy in public places must change as our social environment changes. Having expectations of privacy is particularly relevant when we consider the effects of the fast pace of information technology on our moral norms. For instance, until 1967, when the Supreme Court decided in *Katz v. United States*, that telephone conversations should be private, some considered these conversations to be property of those who owned the telephone equipment. In response to the publication of Robert Bork's videotape rental records in the newspaper, the Video Privacy Act of 1988 reversed the status of video rental records from public to private [27]. More recently, the Violent Crime Control and Law Enforcement Act of 1994 limited access to drivers' records, which were previously regarded as public records ("no-holds barred" [27]). Our use of information technology will continue to reveal new issues of privacy in public places, shifting our societal judgments and affecting our moral norms. Our expectations of privacy will change with such developments.

7 Privacy and the Internet

We maintain that the Internet, in its current state, is a public place. Recall Nissenbaum's separation of place from information [27]. That is, although we may be sitting in a private place (our own homes), when we access the Internet, we are engaging in transactions and exchanging information with other entities (companies, government, educational institutions, etc.) which are definitely not part of our intimate and personal realms. Although we may be in a physically private location, the transactions that take place on the Web strongly parallel those transactions that occur every day in our public lives. If the Internet is a public place, what then should be our reasonable expectations of privacy? Currently, reasonable expectations of privacy on the Internet are neither formally rooted nor well developed. Because Internet technology facilitates the manipulation and collection of information, losses of privacy occur continually, and unfortunately, it is often difficult to determine what constitutes an ethical or unethical collection of data. For example, each time Bob views a page on the Web, that Web site collects several important pieces of information about him: the name of his computer, the time of the request, and the address of the previous Web page he was viewing. As previously discussed,

one way to avoid the unethical collection of personal information is to obtain informed consent. But obtaining informed consent every time Bob moves to a different Web site would be impractical.

It is helpful to compare new Internet situations to more familiar situations, in which our expectations of privacy are better developed. For instance, is making a purchase on the Web similar to purchasing a candy bar from a vending machine? Or is it more similar to purchasing a candy bar from the local supermarket? In order to purchase something on the Web, we must release personal information such as our name, address, and perhaps credit card information. The analogy of the supermarket, then, seems more appropriate than the vending machine. Because we release personal information to a supermarket when we apply for discount cards, we would reasonably expect the supermarket to track our purchases for the purposes of maintaining customer satisfaction and proper inventory. Likewise, we might reasonably expect similar forms of information collection from a Web site.

Such analogies can take us only so far. What should the reasonable expectations of privacy be when we visit virtual galleries or adult sites on the Web? Are and should we be afforded complete anonymity when engaging in online chat groups? It may take time before our moral norms develop in this new Internet context. As we discover new situations in which privacy may be violated on the Internet, we will continue to adjust and reformulate our moral norms.

8 Internet Cookies

Internet cookies provide a test for our claims about the collection and centralization of information. Since early 1996, when an article in the San Jose Mercury brought cookies to public awareness, there has been much concern over cookies' effect on our privacy [29]. Nonetheless, popular sentiment, like the following statement expressed on a public message board on the Internet, demonstrates a failure to distinguish the tool from its use (see section 4.1): I hate cookies. . . They [those who use cookies] may think it's harmless but they are taking something without permission and without payment [30]. Although reasonable expectations of privacy on the Internet are not yet well developed, in this section, we offer an interpretation of a reasonable expectation of privacy with regard to the use of Internet cookies. Once the public has developed a reasonable expectation of privacy for cookies, we can determine what uses of Internet cookies require informed consent. We will show that some uses of cookies are morally permissible while other uses are immoral. Based on our analysis of the collection and centralization of information and our discussions of public places and reasonable expectations, we will identify the conceptual muddles which cookies present and explain how cookies can be used in both morally permissible and unethical ways.

8.1 What are Internet Cookies?

When we visit a Web site, that Web site may give our Web browser a block of text, which is usually a name, value pair. On each subsequent visit to that Web site, our browser sends that specific block of text back to the Web site. Upon receiving that text, the Web site can act in a variety of ways. For instance, it recognizes our browser as a repeat visitor, and may provide us with customized service. It can also change the value of the block of text depending on our behavior at that Web site. Our Web browser remembers this block of text, commonly known as a cookie, by storing it on our hard drive. Not all cookies store information on our hard drives. Transient cookies are stored in the memory of the computer only for the duration of the current web browsing session. For example, Bluestem, the WWW Identification Service at the University of Illinois at Urbana-Champaign uses these transient cookies to store authentication information once users have logged onto the secure system. Mallard, a Web based interactive learning environment developed at the University of Illinois also uses cookies in its authentication system.¹¹

In both cases, the cookies are removed when the user logs out of the system or quits the Web browser. Because these transient cookies seem to pose no apparent ethical problems, the remainder of this paper will focus on the persistent cookies, which are stored on a user's hard drive.

8.2 Argument Against Cookies

Mayer-Schoenberger [31] presents four major reasons why cookies are an invasion of our privacy. In this section we show that the conceptual muddles created by cookies weaken his argument against them.

Cookies are stored on the user's computer without his consent or knowledge. The typical computer user also has no knowledge that cache files, temporary files, and log files are being stored on his computer. In fact, most of these files fill much more space on the hard disk than the small block of text of a cookie. This objection, then, cannot be about cookies, but about the use of computers and technology in general. We might think of a computer as an incomprehensible system. That is, just as we do not necessarily know how our automobile transmission operates, we cannot know about everything which occurs in our computers. In the automobile, we don't need the driver's consent to shift gears. Similarly, we probably do not need consent for a program to store a small file on our hard disk. In fact, it would be counterproductive to our use of computers if we were informed every time an application wanted to change the internal state of our computer. The advantages of Web technology seem to far outweigh the tiny harm of a small block of text set on our hard drive. The cookie is clandestinely and automatically transferred from the user's machine to the Web server. The typical computer user is unaware of much information which is automatically transferred from his computer into the net-

worked world. Each time we visit a Web site, for instance, our computer transmits much information to the Web server including our IP address, the current time, and the previous Web page we were visiting. If our computer is linked to the networked world, it is, without our knowledge, continuously transferring information about its location and existence to other computers. E-mail is not sent directly to the intended receiver, but is automatically transferred through numerous machines of which we have no knowledge. Such automatic transferral of information is essential to our use of current technology. Because cookies allow the Web server to set an expiration date, they violate the "accuracy" and "timeliness" principles in the European Union Directive on the Protection of Personal Data. This argument seems to mistake the tool for its use. In fact, the expiration date option allows the realization of the accuracy and timeliness principles. Of course it is possible for someone to abuse this option thereby violating the principles in the Directive. The cookie, itself, does not violate the principles.¹²

Once the cookie is set, it is freely accessible to Web servers. This argument is technically inaccurate. Only the Web server which set the cookie can access that cookie. Additionally, no other Web server would understand or have use for the cookie except the Web server that set it.

8.3 Morally Permissible Uses of Cookies: Collection of Information

It is important to keep in mind that cookies are merely a tool which is used to collect personal information. As we have discussed previously, the collection of personal information does not necessarily result in a violation of privacy. Imagine that Bob visits his local grocery store. When Bob enters the store, Carol, the store-keeper, immediately recognizes him as a valued repeat customer. She greets Bob with a firm handshake and shows him the new shipment of ripe apples, Bob's favorite fruit, which just arrived. Bob fills his shopping cart with the best fruits and vegetables he can find. He purchases his goods, and takes off down the street, munching on a delicious newly purchased apple. The fact that Carol recognized Bob and remembered what he liked would be considered "doing good business" on the part of Carol. Each time Bob has visited her store, Carol has noticed what types of fruits and vegetables Bob likes to buy. Bob returns to Carol's grocery store because he appreciates the customized service he receives there. When cookies are used for site personalization and online ordering systems, they are an effort to "do good business" on the Web. There are benefits for both the Web site and the Web customer. If Bob likes the ease of service and the personal attention he receives at a Web site, he may return there often. In these situations, the collection of information performed by cookies does nothing more than mimic the memory of Carol. As a repeat customer to Carol's grocery store, Bob has implicitly consented to having Carol remember his preferences.

That is to say, Bob's reasonable expectation of privacy in this scenario is that Carol may recognize him after multiple visits. Likewise, since a large component of the Internet parallels similar purchasing scenarios, it is also within a reasonable expectation of privacy for a Web site to recognize Bob as a repeat visitor. In both cases, the information collected by Carol or the cookie does not result in a violation of privacy because of Bob's reasonable expectations as a repeat customer. Notice that outside the grocery store, Carol knows nothing about Bob. She may not even know Bob's name. Similarly, cookies are useful to a Web site only when we visit that Web site. The Web site may recognize Bob only as someone who has visited before. It does not know his e-mail address, phone number, or home address unless he has explicitly given the Web site such information. Garfinkel [33] describes how cookies may, in fact, be used to protect privacy. Used properly, cookies can eliminate the need for central data banks. Generally, a cookie is a unique number which is used to reference a databank of information stored at the Web site. For instance, Bob may have a cookie which has the value 007 stored on his computer. When Bob visits the Web site, his Web browser sends the number 007 to the site. With that number, the Web site can look up information in its databank and find out that visitor 007 has visited ten times before and likes to read the articles about the apple industry which are available on the site. However, as Garfinkel describes, rather than using the cookie as a unique identifier, the actual preferences might be stored in the cookie itself, eliminating the need for a central databank. For instance, instead of a cookie which has the value 007, the cookie might now consist of:

$$\text{visits} = 10; \text{articles} = \text{apple}$$

Therefore, when Bob leaves, he takes his cookie filled with personal preferences with him, leaving the Web site unable to remember anything about him until he returns.¹³

8.4 Immoral Uses of Cookies: Centralization of Information

Not all uses of cookies are ethical. The use of cookies by the target marketing industry to track our behavior on the Internet is an attempt to centralize personal information. In Section 5, we criticized the centralization of information for taking information out of its intended context and putting it in a new, foreign context. Target marketers' use of cookies is a special case of centralization of information. Their initial purpose in the collection of information is the centralization of information. Target marketers have developed a technique to track us all over the Internet by adding cookies to the advertisement banners on so Web pages. Such uses of cookies do not seem to fit within a reasonable expectation of privacy on the Web. Consider the following scenario: Bob visits the Web site www.dailynews.com to read about the happenings of the day. On that page, he notices an advertisement banner for Jazzy Widgets. The advertisement banner was placed on the www.dailynews.com Web page by a target

marketing company named Banners-R-Us. Unknown to Bob, as he is reading the daily news, a cookie has just been set on his computer by www.bannersrus.com (not www.dailynews.com). After he finishes reading the news, the www.dailynews.com site asks him to register for additional access to the site. Bob happily divulges his name, e-mail, phone, and address because he has enjoyed the Daily News site. However, he does not consent to sending that registration information to Banners-R-Us. Bob completes his visit at www.dailynews.com and decides to visit his favorite online compact disc shopping site, www.columbiahut.com, to purchase some new CDs. At www.columbiahut.com, he notices an advertisement banner for Classical Widgets. Again, the Classical Widgets banner was placed at the Columbia Hut Web site by Banners-R-Us. Again, through the advertisement banner for Classical Widgets and unknown to Bob, the Banners-R-Us cookie previously set at www.dailynews.com is sent to www.bannersrus.com. Banners-R-Us now knows that Bob enjoys reading the news at the Daily News site and purchasing CDs at the Columbia Hut site. If it received any of the registration information Bob gave to the Daily News site, it may even be able to connect his name and e-mail address with his Web browsing behavior [34]. Banners-R-Us will continue to track Bob and gather information about his browsing behavior wherever he bumps into advertisement banners placed by Banners-R-Us on the Web.

The use of cookies to track users as they move from site to site is an unethical invasion of privacy. Such use violates our privacy because it creates an undesirable loss of anonymity and secrecy. No consent has been obtained by target marketers before they collect information about us. Recall from Section 4.3 that the lack of consent does not necessarily render an act unethical. In this cookie case, however, consent seems necessary to legitimize the collection of data. The reason is that such a setting of cookies does not fall within the realm of our current reasonable expectation of privacy for the Web. The cookies set by target marketers differ significantly from those set at a Web site we are actually visiting. When we visit a Web site, it is reasonable to expect that site to collect information about us (recall the example with Carol the storekeeper). When Bob frequents a Web site, he is actively establishing a relationship with that site. With target marketers, however, Bob has no intention nor inclination to establish a relationship with them. For them to obtain information about him without his consent, then, is beyond a reasonable expectation of privacy.

In Section 8.3, we mentioned that a Web site cannot obtain a user's e-mail address, phone number, or home address by setting cookies. According to Cranor [32], however, "multiple Web sites sometimes share access to cookies. A user who reveals personal information to one Web site may unwittingly reveal that information to other sites." Specifically, if Web sites have agreements to share information

with target marketers which have advertisement banners on those sites, information which users may reveal to that particular Web site through registration forms may ultimately be linked to their Web browsing behavior, obtained by target marketers through their use of cookies to centralize information. This is one technique which might be used to generate lists for unsolicited bulk e-mail.

8.5 Cookies and Consent

In section 8.4, we argued that the setting of cookies by target marketers is immoral because such a collection of information is not within a reasonable expectation. We have argued that consent is necessary when a collection of personal information does not comport with a reasonable expectation of privacy.

During the early stages of this paper, the current version of the Netscape Web browser, Netscape Navigator 3.01, had the option to set a "cookie alert" which would warn the user when a site wanted to set a cookie. The alert notification would give the user the option of either accepting or rejecting that cookie. The default setting of Netscape, however, was to accept all cookies unconditionally. Furthermore, there did not exist an option to completely reject all cookies.

Currently, the new Netscape Communicator 4.04 has extended its cookie options. Not only can users set the "cookie alert" option, they also have options to accept all cookies, only cookies that are returned to the originating server, or completely disable cookies. Notice that the option to accept only cookies that are returned to the originating server distinguishes our morally permissible and immoral cookies. The default setting, however, remains the unconditional acceptance of all cookies.¹⁴ In addition, these cookie warnings appear only when a Web site desires to set a cookie on a user's hard drive. Technically, no loss of privacy has occurred at this point. Loss of privacy occurs later, when the cookie is collected and transmitted back to the Web site on subsequent visits. Users must also be informed when such transmission of cookies occur. Mayer-Schoenberger [31] raises the point that the cookie warnings provided by popular Web browsers are cryptic and hard to understand for the typical user. Indeed, these cookie warnings may not provide enough information to be considered informed consent. However, because the default behavior of popular Web browsers is the unconditional acceptance of cookies without notification, only users who are already "cookie savvy" enable the alert option and receive notification. Most users do not even know what a cookie is. Additionally, with the abundant use of cookies on the Internet, cookie warnings soon become a hindrance rather than a help. As Cranor [32] observes, "when such disruptions occur frequently, individuals are unlikely to pay close attention to them."

Cookie preferences should be configured similarly to the security options found on popular Web browsers like Netscape and Internet Explorer. By default, these security

options, unlike cookie options, are enabled. For instance, we are always warned whenever we may be sending unencrypted information to a Web site. To disable this warning, we simply click an option which appears in the warning window itself. In contrast, in order to configure cookie preferences, the user must first know about cookies and then have the patience to find the options in the various browser menus. By default, Web browsers should inform the user about the setting and transmission of cookies. Additionally, the warning windows should then contain the various options to enable, distinguish, and disable all cookies. These options would allow users to configure their cookie preferences "on the y," for example, if the cookies warnings become too annoying.

Another possible solution is to include privacy policy choices (cookie options) in the setup phase of software installation. When installed on a computer, typical software packages must go through a one-time setup phase. During the setup phase, the user is asked to give information which is needed by the software program. Therefore, the setup phase is an ideal place to obtain informed consent from a user. Recall in Section 3.2 that the disclosure of all pertinent information (whether or not the user is interested) is necessary for informed consent. While typical users may not be interested in privacy policies, they must all endure this setup phase in order to use the software. Placing privacy options in this setup phase then forces the user to become informed about the potential privacy losses and violations related to the software. With the recommendations of the previous paragraph, this mechanism would inform users about cookies and give them the ability to change their personal preferences concerning cookies "on the y." Cookies are so prevalent on the Internet that completely disabling them would likely limit the capabilities of the Web browser. We reemphasize that is unnecessary to disable all cookies since typical uses lie within reasonable expectations of privacy. What is important in developing cookie configuration policies, then, is distinguishing morally permissible cookies from immoral cookies.

9 Summary and Conclusions

The Internet provides a new context in which to explore our ideas about privacy. The scale and ease of personal information collection and centralization have caused general concern and confusion regarding our rights to privacy in such an environment. We have offered a framework for evaluating the ethics of the manipulation of information on the World Wide Web. We believe that there are both ethical and unethical ways to collect personal information.

Specifically, collection of information is unethical when such collection lies beyond our reasonable expectation of privacy for the situation. Obtaining informed consent before collecting personal information, however, is a sufficient

means for preventing violations of privacy. While collection of personal information can be ethical, we believe that centralization of personal information is inherently unethical because it undermines the basic function of privacy to create and maintain personal relationships.

Finally, we described the Internet as a public place and offered an interpretation of a reasonable expectation of privacy for certain situations on the Web. Specifically, we explored the issue of Internet cookies and concluded that the use of cookies for online shopping and for customer preferences is morally permissible. In contrast, the use of cookies by target marketers to monitor consumer habits is unethical not only because such collection lies beyond a reasonable expectation of privacy but also because the collected information is unethically centralized. ♦

Acknowledgments:

We thank Willem Bakker II and Helen Nissenbaum for providing direction to this project, James Wallace for the references to the literature on informed consent, Mike Stangel for information on Mallard, and Lorrie Faith Cranor, Lillian Hoddeson, Andrew Pickering, Warden B. Rayward, Ron Szoke, and Marsha Woodbury for suggestions on earlier drafts of this paper.

Notes:

1. Moor asserts that normative privacy is culturally determined: situations which ought to be private should be open to rational and moral argument.
2. Good reason, in this sense, must be a reason grounded on moral principle such as freedom of others, justice, respect for persons, or avoidance of needless pain.
3. Ward had taken her landlord to court because he had failed to deal with the cockroaches and rodents which infested her apartment.
4. In his article "Private Life in Cyberspace", Barlow [19] reflects on the Lotus Marketplace: Households decision. He contrasts today's misuse of information and technology to the restraint exhibited in small towns, where everyone knows information about everyone else, but cares enough to use the information in a respectful manner.
5. We are not saying that the tool itself has no effect on people's behavior. Johnson [16] argues that because computers facilitate the collection of information, people engage in activities which would have not otherwise been possible. Computers, as tools, therefore are an important factor in how people make decisions.
6. An activity is immune if it is not appropriate for unauthorized persons to watch it.
7. As Martin and Schinzingger [15] have observed, most people tend to accept risks which are voluntarily undertaken.
8. The Privacy Act of 1974 forbade the executive branch of the government from sharing information among distinct government program areas. Use of information was limited to "routine use," or a "purpose which is compatible with the purpose for which it was collected." Unfortunately, interpretations of "routine use" were so loose, they left the Privacy Act ineffective. In addition, the Office of Management and Budget, the agency designated to enforce the Privacy Act, essentially refused to uphold the principles of the Act. Both Laudon [20] and Shattuck [21] provide informative sections about the Privacy Act of 1974.
9. Under Gavison's definition of privacy, there may indeed be a loss of privacy. That is to say, because we are the subject of Bob's attention, a degree of our anonymity is lost [8]. However, according to Moor, this would be a loss of natural privacy which is not necessarily a violation of privacy [9].
10. It might be argued that although the centralized information exists, this does not mean that everyone will choose to use it. However, as discussed in 5.2, centralization reduces the cost of obtaining information and makes it more likely that someone will choose to access it. In this sense, it is very much a factor in determining how people act. This argument is similar to the argument Johnson [16] uses to show that tools such as computers are a factor in determining what people do.
11. For more information on Mallard, consult <http://www.cen.uiuc.edu/Mallard>.
12. The European Union Directive on the Protection of Personal Data includes five conditions: (1) personal data must be "processed fairly and lawfully" and only

"collected for a specific, explicit, and legitimate purpose"; (2) no further processing which is incompatible with the original legitimate purpose is permitted; (3) processing must be "adequate, relevant, and not excessive in relation to the purpose" as well as "accurate, and when necessary, kept up to date"; (4) data may be stored for "no longer than necessary for the purposes for which the data was collected"; (5) processing may take place only if the person to whom the personal information refers "has unambiguously given his consent."

13. The contents of the cookies in the above examples are quite simplified. For a detailed description of what cookies actually look like, consult the book by Garfinkel and Spafford [33].
14. While Internet Explorer 4.0 has options to warn, always accept, or disable cookies, it does not distinguish morally permissible cookies from immoral cookies, as Netscape Communicator 4.04 does.

References

- [1] Morin, Richard, "Cruising the Internet...but Warily," in *The Washington Post National Weekly Edition*, June 23, 1997, page 35.
- [2] Moor, James H., "Towards a Theory of Privacy in the Information Age," in *Computers and Society*, September 1997, pp. 27-32.
- [3] Warren, Samuel D. and Louis D. Brandeis, "The Right to Privacy," in Schoeman, Ferdinand David, ed., *Philosophical Dimensions of Privacy: An Anthology*, Cambridge, Massachusetts: Cambridge University Press, 1984, pp. 75-103.
- [4] Fried, Charles, "Privacy," in Schoeman, Ferdinand David, ed., *Philosophical Dimensions of Privacy: An Anthology*, Cambridge, Massachusetts: Cambridge University Press 1984, pp. 203-222.
- [5] Westin, Alan F. *Privacy and Freedom*, New York, Atheneum, 1967.
- [6] Beardsley, Elizabeth, "Privacy: Autonomy and selective disclosure," in Pennock, J.R. and J.W. Chapman, eds., *Nomos XIII: Privacy*, New York: Atherton Press, 1971.
- [7] Parent, W.A. "Privacy, Morality, and the Law," in *Philosophy and Public Affairs*, vol. 12, no. 4, Fall 1983, pp. 269-288.
- [8] Gavison, Ruth, "Privacy and the Limits of Law," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 332-350.
- [9] Moor, James H., "The Ethics of Privacy Protection," in *Library Trends*, vol. 39, nos. 1 and 2, Summer/Fall 1990, pp. 69-82.
- [10] Rachels, James. "Why Privacy is Important," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 351-357.
- [11] Benn, Stanley I. "Privacy, freedom, and respect for persons," in Schoeman, Ferdinand David, ed., *Philosophical Dimensions of Privacy: An Anthology*, Cambridge, Massachusetts: Cambridge University Press, 1984, pp. 223-234.
- [12] Faden, Ruth R. and Tom L. Beauchamp. *A History and Theory of Informed Consent*. New York: Oxford University Press, 1986.
- [13] Gorovitz, Samuel. "Informed Consent and Patient Autonomy," in Callahan, Joan C., ed., *Ethical Issues in Professional Life*. New York: Oxford University Press, 1988, pp. 182-187.
- [14] Lidz, Charles, Alan Meisel, Eviatar Zerubavel, Mary Carter, Regina M. Sestak, and Loren H. Roth. *Informed Consent: A Study in Decisionmaking in Psychiatry*. New York: The Guilford Press, 1984.
- [15] Martin, Mike W. and Roland Schinzinger. *Ethics in Engineering*. McGraw-Hill Publishing Company, 1989.
- [16] Johnson, Deborah G. *Computer Ethics*, 2nd ed., Upper Saddle River, New Jersey, Prentice Hall, 1994.
- [17] Forester, Tom and Perry Morrison. *Computer Ethics: Cautionary Tales and Ethical Dilemmas in Computing*, 2nd ed., Cambridge, Massachusetts: The MIT Press, 1994.
- [18] Quittner, Joshua. "Invasion of Privacy" in *Time*, August 25, 1997, pp. 28-35.
- [19] Barlow, John P. "Private Life in Cyberspace," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computer Ethics & Social Values*, Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 310-313.
- [20] Laudon, Kenneth C. *Dossier Society: Value Choices in the Design of National Information Systems*, New York: Columbia University Press, 1986. 30
- [21] Sharrock, John. "Computer Matching is a Serious Threat to Individual Rights," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 305-309.
- [22] Culnan, M.J. and H.J. Smith. "Lotus Marketplace: Households... Managing Information Privacy Concerns," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 269-277.
- [23] "Where is our Data?" in *Secure Computing: The Magazine for the Protection of Information*, August 1997, pp. 18-22.
- [24] Hunter, Larry, "Public Image," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 293-298.
- [25] Parker, Donn B., *Ethical Conflicts in Computer Science and Technology*, Arlington, Virginia: AFIPS Press, 1979.
- [26] Ware, Willis H., "The New Faces of Privacy," in *The Information Society*, vol. 9, 1993, pp. 195-211.
- [27] Nissenbaum, Helen. "Toward an Approach to Privacy in Public: Challenges of Information Technology," in *Ethics and Behavior*, vol. 7, no. 3, 1997, pp. 207-219.
- [28] Gentry, Tom. "Redefining Privacy," in *Harvard Civil Rights- Civil Liberties Law Review*, vol. 12, no. 2, 1977.
- [29] Malcolm's Guide to Persistent Cookies <http://www.emf.net/~mal/cookiesinfo.html>
- [30] Andy's HTTP Cookie Notes: <http://www.illuminatus.com/cookie.cgi>
- [31] Mayer-Schoenberger, Viktor. "The Internet and Privacy Legislation: Cookies for a Treat?" in *West Virginia Journal of Law and Technology*, vol. 1, issue 1, March 17, 1997, at <http://www.wvijolt.wvu.edu/wvijolt/current/issue1/articles/mayer/mayer.htm>.
- [32] Cranor, Lorrie Faith. "The Role of Technology in Self-Regulatory Privacy Regimes," prepared for the National Telecommunications and Information Administration, December 1996, at <http://www.research.att.com/~lorrie/pubs/NTIA.html>. 31
- [33] Garfinkel, Simon with Gene Spafford. *Web Security and Commerce*. O'Reilly & Associates, Inc., 1997.
- [34] Privacy section at <http://www.cookiecentral.com/javacook2.htm>.
- [35] Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995.
- [36] Nissenbaum, Helen, "Can We Protect Privacy in Public?" in *ACM SIGCAS Conference in Computer Ethics*. June 1997, Rotterdam, The Netherlands.
- [37] Moor, James, H., "What is Computer Ethics," in Johnson, Deborah G. and Helen Nissenbaum, eds., *Computers, Ethics, & Social Values*. Upper Saddle River, New Jersey: Prentice Hall, 1995, pp. 7-14.

Permission to make digital/hard copy of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication and its date appear, and notice is given that copying is by permission of ACM, Inc. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POLICY '98, Washington, DC © 1998 ACM 1-58113-038-4/98/0500 \$5.00

XI. Related Proceedings Appendix

No decisions made in related appeal for U.S. patent application no. 10/995,295.